



US008745372B2

(12) **United States Patent**
Orsini et al.

(10) **Patent No.:** **US 8,745,372 B2**

(45) **Date of Patent:** **Jun. 3, 2014**

(54) **SYSTEMS AND METHODS FOR SECURING DATA IN MOTION**

USPC 380/37, 42, 255; 705/76; 713/168, 186, 713/150, 156, 175, 151, 155; 726/15
See application file for complete search history.

(75) Inventors: **Rick L. Orsini**, Flower Mound, TX (US); **Mark S. O'Hare**, Coto De Caza, CA (US); **Stephen C. Bono**, Baltimore, MD (US); **Gabriel D. Landau**, Baltimore, MD (US); **Seth James Nielson**, Lutherville, MD (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,453,074 A 6/1984 Weinstein
4,924,513 A 5/1990 Herbison et al.

(Continued)

FOREIGN PATENT DOCUMENTS

AU 2004248616 6/2009
EP 346180 12/1989

(Continued)

OTHER PUBLICATIONS

(73) Assignee: **Security First Corp.**, Rancho Santa Margarita, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 232 days.

(21) Appl. No.: **12/953,877**

(22) Filed: **Nov. 24, 2010**

(65) **Prior Publication Data**

US 2011/0202755 A1 Aug. 18, 2011

Related U.S. Application Data

(60) Provisional application No. 61/264,464, filed on Nov. 25, 2009.

(51) **Int. Cl.**

H04L 29/06 (2006.01)
G06F 15/16 (2006.01)
H04L 9/00 (2006.01)
H04L 12/46 (2006.01)

(52) **U.S. Cl.**

CPC **H04L 9/00** (2013.01); **H04L 63/0272** (2013.01); **H04L 63/0823** (2013.01); **H04L 12/4641** (2013.01)
USPC **713/151**; 713/155; 713/156; 713/168; 726/15

(58) **Field of Classification Search**

CPC H04L 9/00; H04L 9/006; H04L 9/0827; H04L 12/4633; H04L 12/4641; H04L 29/06612; H04L 63/0272; H04L 63/0823; H04L 2012/5621

Pattam, "Enhancing Security in 802.11 and 802.1X Networks with Intrusion Detection", Jan. 20, 2006, University of New Orleans Theses and Dissertations, Paper 1034, 111 pages.*

(Continued)

Primary Examiner — Andrew Nalven

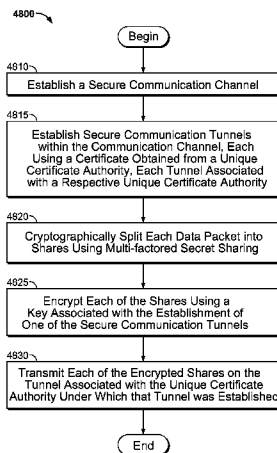
Assistant Examiner — Walter Malinowski

(74) *Attorney, Agent, or Firm* — Ropes & Gray LLP

(57) **ABSTRACT**

Systems and methods are provided for distributing trust among a set of certificate authorities. One approach provides methods and systems in which the secure data parser is used to distribute trust in a set of certificate authorities during initial negotiation of a connection between two devices. Another approach provides methods and systems in which the secure data parser is used to disperse packets of data into shares. A set of tunnels is established within a communication channel using a set of certificate authorities, keys developed during the establishment of the tunnels are used to encrypt shares of data, and the shares of data are transmitted through each of the tunnels.

37 Claims, 53 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

4,932,057 A	6/1990	Kolbert	6,386,451 B1	5/2002	Sehr
5,010,572 A	4/1991	Bathrick et al.	6,411,716 B1	6/2002	Brickell
5,016,274 A	5/1991	Micali et al.	6,415,373 B1	7/2002	Peters et al.
5,051,745 A	9/1991	Katz	6,424,718 B1	7/2002	Holloway
5,150,407 A	9/1992	Chan	6,438,690 B1	8/2002	Patel et al.
5,163,096 A	11/1992	Clark et al.	6,446,204 B1	9/2002	Pang et al.
5,268,963 A	12/1993	Monroe et al.	6,449,730 B2	9/2002	Mann et al.
5,375,244 A	12/1994	McNair	6,453,416 B1	9/2002	Epstein
5,386,104 A	1/1995	Sime	6,483,921 B1	11/2002	Harkins
5,450,099 A	9/1995	Stephenson et al.	6,505,216 B1	1/2003	Schutzman et al.
5,485,474 A	1/1996	Rabin	6,553,493 B1	4/2003	Okumura et al.
5,524,073 A	6/1996	Stambler	6,615,347 B1	9/2003	de Silva et al.
5,603,003 A	2/1997	Akizawa et al.	6,625,734 B1	9/2003	Marvit et al.
5,615,269 A	3/1997	Micali	6,631,201 B1	10/2003	Dickinson et al.
5,623,549 A *	4/1997	Ritter 380/37	6,684,330 B1	1/2004	Wack et al.
5,642,508 A	6/1997	Miyazawa	6,687,375 B1	2/2004	Matyas, Jr. et al.
5,666,414 A	9/1997	Micali	6,701,303 B1 *	3/2004	Dunn et al. 705/75
5,666,416 A	9/1997	Micali	6,731,755 B1	5/2004	Cocks
5,666,514 A	9/1997	Cheriton	6,785,768 B2	8/2004	Peters et al.
5,682,425 A	10/1997	Enari	6,789,198 B1	9/2004	Chan
5,703,907 A	12/1997	James	6,819,766 B1	11/2004	Weidong
5,717,758 A	2/1998	Micall	6,852,988 B2	2/2005	Li
5,748,735 A	5/1998	Ganesan	6,856,383 B1	2/2005	Vachris et al.
5,761,306 A	6/1998	Lewis	6,915,436 B1 *	7/2005	Booth et al. 726/3
5,768,382 A	6/1998	Schneider et al.	6,952,684 B2	10/2005	Toshikage et al.
5,768,519 A	6/1998	Swift et al.	6,959,383 B1	10/2005	Terada et al.
5,790,677 A	8/1998	Fox et al.	6,965,993 B2	11/2005	Baker
5,823,948 A	10/1998	Ross, Jr. et al.	6,978,367 B1	12/2005	Hind et al.
5,903,652 A	5/1999	Mital	7,003,531 B2	2/2006	Holenstein et al.
5,903,882 A	5/1999	Asay et al.	7,050,580 B1	5/2006	Ferre Herrero
5,910,987 A	6/1999	Ginter et al.	7,050,583 B2	5/2006	Montgomery
5,915,019 A	6/1999	Ginter et al.	7,069,451 B1	6/2006	Ginter et al.
5,917,913 A	6/1999	Wang	7,072,917 B2	7/2006	Wong et al.
5,937,066 A	8/1999	Gennaro et al.	7,085,854 B2 *	8/2006	Keane et al. 709/250
5,940,507 A	8/1999	Cane et al.	7,096,494 B1	8/2006	Chen
5,960,083 A	9/1999	Micali	7,107,385 B2	9/2006	Rajan et al.
5,966,444 A	10/1999	Yuan et al.	7,117,365 B1	10/2006	Rump et al.
5,966,448 A	10/1999	Namba et al.	7,133,845 B1	11/2006	Ginter et al.
5,974,144 A	10/1999	Brandman	7,143,289 B2	11/2006	Denning et al.
5,982,322 A	11/1999	Bickley et al.	7,171,685 B2 *	1/2007	Batra et al. 726/14
5,987,232 A	11/1999	Tabuki	7,174,385 B2	2/2007	Li
5,991,414 A	11/1999	Garay et al.	7,187,771 B1	3/2007	Dickinson et al.
5,995,630 A	11/1999	Borza	7,203,871 B2	4/2007	Turner et al.
6,009,177 A	12/1999	Sudia	7,218,736 B1	5/2007	Nishimura et al.
6,023,508 A	2/2000	Bombard et al.	7,225,158 B2	5/2007	Toshikage et al.
6,026,163 A	2/2000	Micali	7,277,958 B2	10/2007	Chung et al.
6,061,790 A	5/2000	Bodnar	7,302,583 B2	11/2007	Forrest
6,073,237 A	6/2000	Ellison	7,304,990 B2	12/2007	Rajwan
6,092,201 A	7/2000	Turnbull et al.	7,337,320 B2	2/2008	Tada et al.
6,094,485 A	7/2000	Weinstein et al.	7,349,539 B2	3/2008	Wee et al.
6,118,874 A	9/2000	Okamoto et al.	7,391,865 B2	6/2008	Orsini et al.
6,134,550 A	10/2000	Van Oorschot et al.	7,412,462 B2	8/2008	Margolus et al.
6,192,472 B1	2/2001	Garay et al.	7,428,751 B2	9/2008	Oom Temudo de Castro et al.
6,229,894 B1	5/2001	Van Oorschot et al.	7,428,754 B2 *	9/2008	Neumann et al. 726/15
6,240,183 B1	5/2001	Marchant	7,444,421 B2	10/2008	Katayama
6,240,187 B1	5/2001	Lewis	7,472,105 B2	12/2008	Staddon et al.
6,256,737 B1	7/2001	Bianco et al.	7,535,905 B2	5/2009	Narayanan et al.
6,260,125 B1	7/2001	McDowell	7,546,427 B2	6/2009	Gladwin et al.
6,266,413 B1	7/2001	Shefi	7,596,570 B1	9/2009	Emigh et al.
6,269,432 B1	7/2001	Smith	7,627,125 B2	12/2009	Lumsden et al.
6,289,455 B1	9/2001	Kocher et al.	7,649,992 B2	1/2010	Raju et al.
6,289,509 B1	9/2001	Kryloff	7,693,992 B2	4/2010	Watson
6,292,568 B1	9/2001	Akins, III et al.	7,721,150 B2	5/2010	Belyakov et al.
6,292,782 B1	9/2001	Weideman	7,752,482 B2	7/2010	Gredone et al.
6,301,659 B1	10/2001	Micali	7,801,781 B2 *	9/2010	Olin et al. 705/34
6,307,940 B1	10/2001	Yamamoto et al.	7,865,741 B1 *	1/2011	Wood et al. 713/193
6,314,409 B2	11/2001	Schneck et al.	7,869,597 B2	1/2011	Nakai et al.
6,324,650 B1	11/2001	Ogilvie	7,921,288 B1	4/2011	Hildebrand
6,336,186 B1	1/2002	Dyksterhouse et al.	8,151,333 B2	4/2012	Zhu et al.
6,345,101 B1	2/2002	Shukla	8,347,084 B2 *	1/2013	Tavernier et al. 713/162
6,345,314 B1	2/2002	Cole et al.	8,397,288 B2 *	3/2013	Melvin et al. 726/15
6,356,941 B1	3/2002	Cohen	2001/0001876 A1	5/2001	Morgan et al.
6,363,425 B1	3/2002	Hook et al.	2001/0051902 A1	12/2001	Messner
6,363,481 B1	3/2002	Hardjono	2002/0010679 A1	1/2002	Felsher
6,385,318 B1	5/2002	Oishi	2002/0016926 A1 *	2/2002	Nguyen et al. 713/201
			2002/0023210 A1 *	2/2002	Tuomenoksa et al. 713/161
			2002/0032663 A1	3/2002	Messner
			2002/0046359 A1	4/2002	Boden
			2002/0071566 A1	6/2002	Kurn

(56)

References Cited

U.S. PATENT DOCUMENTS

2002/0075844 A1* 6/2002 Hagen 370/351
 2002/0080888 A1* 6/2002 Shu et al. 375/295
 2002/0087866 A1 7/2002 Berson et al.
 2002/0091640 A1 7/2002 Gupta
 2002/0129235 A1 9/2002 Okamoto et al.
 2002/0129245 A1 9/2002 Cassagnol et al.
 2002/0157007 A1 10/2002 Sashihara
 2002/0162047 A1 10/2002 Peters et al.
 2002/0172358 A1 11/2002 Hurich
 2002/0178355 A1* 11/2002 D'Sa et al. 713/156
 2002/0178361 A1* 11/2002 Genty et al. 713/175
 2002/0184444 A1* 12/2002 Shandony 709/203
 2003/0005094 A1* 1/2003 Yuan et al. 709/221
 2003/0023958 A1 1/2003 Patel et al.
 2003/0051054 A1 3/2003 Redlich et al.
 2003/0051159 A1 3/2003 McCown et al.
 2003/0058274 A1 3/2003 Hill et al.
 2003/0061481 A1* 3/2003 Levine et al. 713/163
 2003/0070077 A1 4/2003 Redlich et al.
 2003/0084020 A1 5/2003 Shu
 2003/0084290 A1 5/2003 Murty et al.
 2003/0097564 A1 5/2003 Tewari et al.
 2003/0167408 A1 9/2003 Fitzpatrick et al.
 2004/0093431 A1* 5/2004 Genty et al. 709/245
 2004/0111608 A1 6/2004 Oom Temudo de Castro et al.
 2004/0122863 A1 6/2004 Sidman
 2004/0122960 A1* 6/2004 Hall et al. 709/229
 2004/0220879 A1 11/2004 Hughes
 2004/0225895 A1* 11/2004 Mukherjee et al. 713/201
 2004/0268148 A1* 12/2004 Karjala et al. 713/201
 2005/0050004 A1 3/2005 Sheu et al.
 2005/0132060 A1* 6/2005 Mo et al. 709/227
 2005/0132221 A1* 6/2005 Marcjan 713/201
 2005/0160290 A1* 7/2005 Moon et al. 713/201
 2005/0198306 A1* 9/2005 Palojarvi et al. 709/227
 2006/0046728 A1* 3/2006 Jung et al. 455/445
 2006/0075222 A1* 4/2006 Moloney et al. 713/156
 2006/0165060 A1* 7/2006 Dua 370/352
 2006/0177061 A1* 8/2006 Orsini et al. 380/268
 2006/0184764 A1 8/2006 Osaki
 2006/0232826 A1 10/2006 Bar-El
 2006/0282681 A1 12/2006 Scheidt et al.
 2006/0294378 A1 12/2006 Lumsden et al.
 2007/0006015 A1 1/2007 Rao et al.
 2007/0033644 A1* 2/2007 Thomas et al. 726/10
 2007/0064704 A1* 3/2007 Balay et al. 370/392
 2007/0071234 A1 3/2007 Lagrange et al.
 2007/0079082 A1* 4/2007 Gladwin et al. 711/154
 2007/0097885 A1* 5/2007 Traversat et al. 370/254
 2007/0124584 A1 5/2007 Gupta
 2007/0157025 A1 7/2007 Sastry et al.
 2007/0160198 A1* 7/2007 Orsini et al. 380/28
 2007/0162744 A1* 7/2007 Hoshino et al. 713/156
 2007/0177894 A1* 8/2007 Nakano et al. 399/121
 2007/0192586 A1 8/2007 McNeely
 2007/0281664 A1* 12/2007 Kaneko et al. 455/410
 2008/0016334 A1* 1/2008 Kurapati et al. 713/155
 2008/0037557 A1* 2/2008 Fujita et al. 370/395.53
 2008/0046752 A1* 2/2008 Berger et al. 713/186
 2008/0046763 A1* 2/2008 Teramoto et al. 713/194
 2008/0072035 A1* 3/2008 Johnson et al. 713/153
 2008/0082831 A1* 4/2008 Fujimaki 713/178
 2008/0108349 A1* 5/2008 Ihattula 455/435.1
 2008/0126614 A1 5/2008 Ooi et al.
 2008/0137855 A1* 6/2008 Enomoto et al. 380/255
 2008/0137857 A1 6/2008 Bellare et al.
 2008/0155252 A1* 6/2008 Nambiar 713/153
 2008/0170693 A1 7/2008 Spies et al.
 2008/0183992 A1* 7/2008 Martin et al. 711/162
 2008/0199003 A1 8/2008 Hennessey et al.
 2008/0240441 A1 10/2008 Kawakami
 2008/0244277 A1 10/2008 Orsini et al.
 2008/0270603 A1* 10/2008 Berger et al. 709/224
 2009/0016357 A1* 1/2009 Blevins et al. 370/395.53
 2009/0077379 A1 3/2009 Geyzel et al.

2009/0089879 A1 4/2009 Wang et al.
 2009/0092252 A1* 4/2009 Noll et al. 380/277
 2009/0097657 A1* 4/2009 Scheidt et al. 380/277
 2009/0097661 A1* 4/2009 Orsini et al. 380/279
 2009/0113203 A1* 4/2009 Tsuge et al. 713/151
 2009/0113424 A1 4/2009 Chen et al.
 2009/0177894 A1* 7/2009 Orsini et al. 713/193
 2009/0209232 A1* 8/2009 Cha et al. 455/411
 2009/0254750 A1 10/2009 Bono et al.
 2009/0292919 A1* 11/2009 England 713/168
 2009/0300351 A1 12/2009 Lei et al.
 2009/0304004 A1 12/2009 Huynh Van et al.
 2009/0305900 A1 12/2009 Belouchi et al.
 2010/0005028 A1 1/2010 Hartley et al.
 2010/0054474 A1* 3/2010 Schneider 380/277
 2010/0077252 A1 3/2010 Siewert et al.
 2010/0125730 A1 5/2010 Dodgson et al.
 2010/0150341 A1 6/2010 Dodgson et al.
 2010/0153670 A1 6/2010 Dodgson et al.
 2010/0153703 A1 6/2010 Dodgson et al.
 2010/0154053 A1 6/2010 Dodgson et al.
 2010/0223613 A1 9/2010 Schneider
 2010/0287374 A1* 11/2010 Roy et al. 713/171
 2010/0299313 A1 11/2010 Orsini et al.
 2010/0332617 A1 12/2010 Goodwin et al.
 2011/0085667 A1* 4/2011 Berrios et al. 380/282
 2011/0093113 A1* 4/2011 Sager et al. 705/318
 2011/0179271 A1 7/2011 Orsini et al.
 2011/0179287 A1 7/2011 Orsini et al.
 2011/0296104 A1 12/2011 Noda et al.
 2012/0266231 A1* 10/2012 Spiers et al. 726/12
 2013/0042105 A1 2/2013 Orsini et al.

FOREIGN PATENT DOCUMENTS

EP 354774 2/1990
 EP 0485090 A2 5/1992
 EP 636259 2/1995
 EP 0695997 A2 2/1996
 EP 793367 9/1997
 EP 0821504 1/1998
 EP 0862301 9/1998
 EP 1011222 6/2000
 EP 1239384 A2 9/2002
 GB 2237670 A 5/1991
 JP 04297157 A 10/1992
 JP 2006174152 A * 6/2006
 RU 2124814 1/1999
 WO WO-98/4709 A1 10/1998
 WO WO-99/19845 A1 4/1999
 WO WO-99/46720 A1 9/1999
 WO WO-99/65207 A1 12/1999
 WO WO-00/36786 A1 6/2000
 WO WO-00/76118 A1 12/2000
 WO WO-00/79367 A1 12/2000
 WO WO-01/22201 A1 3/2001
 WO WO-01/22319 A1 3/2001
 WO WO-01/22322 A2 3/2001
 WO WO-01/22650 A2 3/2001
 WO WO-01/22651 A2 3/2001
 WO WO-02/21283 A1 3/2002
 WO WO-02/21761 A2 3/2002
 WO WO 02062032 A2 * 8/2002
 WO WO-2004/111791 A2 12/2004
 WO WO-2006/047694 A1 5/2006
 WO WO-2008/054406 A2 5/2008
 WO WO-2008/070167 A1 6/2008
 WO WO-2008/127309 A2 10/2008
 WO WO-2008/142440 A1 11/2008
 WO WO-2009/035674 A1 3/2009
 WO WO-2009/089015 A1 7/2009
 WO WO-2009/105280 A2 8/2009
 WO WO-2010/057181 A2 5/2010
 WO WO-2010/135412 A2 11/2010

(56)

References Cited

FOREIGN PATENT DOCUMENTS

WO WO-2011/068738 A2 6/2011
 WO WO-2011/123692 A2 10/2011
 WO WO-2011/123699 A2 10/2011

OTHER PUBLICATIONS

- Dierks, "The Transport Layer Security Protocol", RFC 4346, Version 1.1, Apr. 2006, 86 pages.*
- Kent, "Security Architecture for the Internet Protocol", RFC 2401, Nov. 1998, 67 pages.*
- Grant, Lawrence, and Brian Fleming. "Secret Sharing and Splitting." White Paper) Notre Dame, Indiana (2002).*
- Sengar, "MTPSec: Customizable Secure MTP3 Tunnels in the SS7 Network", Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium, 2005, 8 pages.*
- Mondal, "Mobile IP", Chapter 7, Kluwer Academic/ Plenum Publishers, New York, 2003, pp. 173-202.*
- Shaw, "Handbook on Electronic Commerce", Conway, Chapter 31, Springer-Verlag, Berlin Heidelberg, 2000, pp. 649-690.*
- Hancock, "Network Security", "IPV6 Security Enhancements Still Not Everything You Need", Elsevier Science Ltd., Jun. 1998, pp. 8-10.*
- Nacht, "The Spectrum of Modern Firewalls", Computers & Security, 16 (1997), pp. 54-56.*
- Sengar, "MTPSec: Customizable Secure MTP3 Tunnels in the SS7 Network", Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IP'05), 8 pages.*
- Rosen, "BGP/MPLS VPNs", Network Working Group, RFC 2547, Mar. 1999, 25 pages.*
- "Decru Unveils Security Appliances for Storage Networks; Decru DataFort (TM) Security Alliances Protect SAN and NAS Environments with Wire-Speed Encryption and Transparent Deployment", PR Newswire (PR Newswire Association, Inc.), Oct. 14, 2002.
- "Lancope Announces Stealthwatch 3.0 for Enhanced Enterprise-Wide Security and Improved Manageability", Business Wire (Newswire Association, Inc.), Apr. 14, 2003.
- "Tactilesense TM White Paper—A Breakthrough in Fingerprint Authentication", Ethentica, Inc. By Security First Corporation, Jan. 2003.
- "Trustengine TM White Paper—Enthentication Services, Secure Storage and Authentication Solutions", Ethentica, Inc. by Security First Corporation, Jun. 2002.
- Barlas, "RSA's Security Showcase", Line56.com—The E-Business Executive Daily, Apr. 15, 2003.
- Brainard, "A New Two-Server Approach for Authentication with Short Secrets" RSA Laboratories, Bedford, MA (13 pages).
- Cachin, "On-Line secret Sharing," Cryptography and Coding. IMA Conference, Proceedings, XX, XX, Dec. 18, 1995, pp. 190-198, XP002137681.
- Chan et. al., "Distributed Server Networks for Secure Multicast", Globcom'01: IEEE Global Telecommunications Conference (IEEE, Piscataway, NJ) 3:1974-1978 (2001).
- Chan et. al., "Distributed Servers Approach for Large-Scale Multicast", IEEE Journal on Selected Areas in Communications (IEEE, Piscataway, NJ) Oct. 2002, 20(8):1500-1510.
- Crescenzo et al., "Non-Interactive and Non-Malleable Commitment," Proceedings of the 30th Annual ACM Symposium on Theory of Computing, Dallas, TX, May 23-26, 1998, [Proceedings of the 30th Annual ACM Symposium on Theory of Computing], New York, NY: ACM, US, pp. 141-150; XP000970902; ISBN: 978-0-89791-962-3.
- Damgard et. al., "Non-interactive and Reusable Non-malleable Commitment Schemes," ACM STOC '03; pp. 427-428; Jun. 9-11, 2003.
- Doyle, "RSA Splits Data to Stop Hackers", vnunet.com, Apr. 16, 2003.
- Fisher, "RSA Looks to Lock Down Personal Data", eWeek—Enterprise News & Reviews, Apr. 14, 2003.
- Garay et. al., "Secure distributed storage and retrieval," Theoretical Comput. Sci., 243(1-2):363-389, Jul. 2000.
- Gibson, "Opinion", eWeek—Enterprise News & Reviews, Apr. 14, 2003.
- Grant et. al., "Secret Sharing and Splitting", (White Paper) Notre Dame, Indiana, Dec. 16, 2002.
- Hunter, "Simplifying PKI Usage Through a Client-Server Architecture and Dynamic Propagation of Certificate Paths and Repository Addresses", Proceedings 13th International Workshop on Database and Expert Systems Applications (IEEE Computer Soc., Los Alamitos, CA), Sep. 2-6, 2002, p. 505-510.
- International Search Report—International Application No. PCT/US/06/45066, dated Jul. 17, 2008.
- International Search Report and Written Opinion dated Dec. 14, 2010 in International Application No. PCT/US2010/035377.
- International Search Report and Written Opinion dated Sep. 8, 2009 in International Application No. PCT/US2009/001158.
- International Search Report dated Dec. 16, 2008, International Application No. PCT/US07/023626.
- International Search Report dated Mar. 10, 2009, International Application No. PCT/US09/00083.
- International Search Report dated Nov. 21, 2008, International Application No. PCT/US08/10677.
- Krawczyk, "Distributed Fingerprints and Secure Information Dispersal," 12th ACM, Symposium on Principles on Distributed Computing, Ithaca, NY, ACM 0-89191-613-1/93/0008/0207, 1993, pp. 207-218.
- Krawczyk, "Secret sharing made Short" IBM T.J. Watson Research Center, [Online] 1998, retrieved from the Internet: URL:<http://www.cs.cornell.edu/courses/cs754/2001fa/secretshort.pdf> [retrieved on Nov. 24, 2008].
- Loutre et. al. "An EAP-BT Smartcard for Authentication in the Next Generation of Wireless Communications", Conference on Network Control and Engineering for QoS Security and Mobility (Kluwer Academic Publishers, Norwell, MA) Oct. 23-25, 2002, pp. 1-4-114.
- Mayer et. al., "Generalized Secret Sharing and Group-Key Distribution Using Short Keys," Compression and Complexity of Sequences 1997, Proceedings Salerno, Italy, Jun. 11-13, 1997, Los Alamitos, CA, USA, IEEE Comput. Soc, US, Jun. 11, 1997, pp. 30-44, XP010274905, ISBN: 978-0/8186-8132-5.
- McNamara, "Strong Crypto Freeware" (Secret Sharer Version 1.0) Jul. 11, 1995.
- Menezes, "Handbook of Applied Cryptography" 1997, CRC Press LLC, XP001525007.
- Mitchell, "Making Serial Number Based Authentication Robust Against Loss of State", Operating Systems Review, 34:3, pp. 1-5, Nov. 24, 1999.
- Myers et. al., "A secure, publisher-centric Web caching infrastructure" In: Infocom 2001 Proceedings. IEEE Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies [online], vol. 3 p. 1235-1243. Published Apr. 22, 2001. [retrieved on Jul. 8, 2008]. Retrieved from the internet <URL: <http://people.ischool.berkeley.edu/~chuang/pubs/gemini.pdf>>.
- Nightingale, The New Secret-Splitting Technology from RSA . . . NGBK DS 0403 <http://developer.rsasecurity.com/labs/nightingale/developer.rsasecurity.com/labs/nightingale/files/nightingale-brochure.pdf>.
- Rabin, "Efficient Dispersal of Information for Security, Load Balancing and Fault Tolerance," Journal of the Association for Computing Machinery, vol. 36, No. 2, pp. 335-348, Apr. 1989.
- RSA SureFile: Software Powered by PKZIP . . . BSSF DS 0103 Authorized Reseller: Technical Specifications Platforms Microsoft® Windows® 98 Second Edition ME NT 4.0 Workstation SP6A 2000 Professional SP2 . . . www.rsasecurity.com/products/bsafe/datasheets/BSSF_DS_103.pdf.
- Savage, "RSA Unveils Nightingale Technology", CRN.com, Apr. 14, 2003.
- Shamir, "How to Share a Secret," Communications of the Association for Computing Machinery, ACM, New York, NY, US, vol. 22, No. 11, Nov. 1, 1979, pp. 1-4, XP002241399; ISSN: 0001-0782.
- Shin et. al., "Design a Working Model of Secure Data Transfer Using a Data Mart", Proceedings of the ISCA 14th International Conference Computer Applications in Industry and Engineering (ISCA, Cary, NC) Nov. 27-29, 2001, pp. 66-69.

(56)

References Cited

OTHER PUBLICATIONS

- Vijayan, "RSA unveils Management, Encryption Products", Computerworld, Apr. 15, 2003.
- Waldman et al., "Publius: A robust, tamper-evident, censorship-resistant web publishing system," Proceedings of the 9th Usenix Security Symposium, Aug. 2000.
- Waters, "RSA Integrates Id Management; discloses 'Nightingale'", ADTmag.com, Apr. 21, 2003.
- Baughner et al., "MSEC Group Key Management Architecture draft-ietf-msec-gkmarch-06.txt; draft-ietf-msec-gkmarch-06.txt," vol. MSEC, No. 6, Sep. 8, 2003.
- Butler et al., "Privacy Preserving Web-Based Email," Info. Systems Security Lecture Notes in Computer Science; Springer, Berlin, Germany, pp. 116-131.
- Declene et al., "Secure group communications for wireless networks," Milcom 2001 Proceedings, Communications for Network-Centric Operations: Creating the Information Force; IEEE, US, vol. 1, Oct. 28, 2001, pp. 113-117.
- Easter et al. "S/390 parallel enterprise server CMOS cryptographic coprocessor," IBM Journal of Research and Development, International Business Machines Corporation, New York, NY, US, vol. 43, No. 5, Jan. 1, 1999, pp. 761-776, XP002335589, ISSN: 0018-8646.
- Ertaul et al., "ECC Based Threshold Cryptography for Secure Data Forwarding and Secure Key Exchange in MANET (I)", Networking 2005. Networking Technologies, Services and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems; vol. 3462/2005, pp. 102-113.
- Ganger et al., "PISIS: A Distributed Framework for Perpetually Available and Secure Information Systems, Final Technical rept. Jun. 1999-Dec. 2003," (Jul. 1, 2005), pp. 1-302, XP55011444, Retrieved from the Internet: URL: <http://www.dtic/mil/cgi-bin/GetTRDoc?AD=ADA436245&Location=U2&doc=GetTRDoc.pdf> (retrieved on Nov. 7, 2011).
- Ganger et al., "Survivable storage systems," DARPA Information Survivability Conference & Exposition II, 2001. Disc Ex '01. Proc. Jun. 12-14, 2001, Piscataway, NJ, USA, IEEE, vol. 2, pp. 184-195, XP010548746.
- Hand et al., Spread Spectrum Storage with Mnemosyne, 2003, Retrieved from the Internet <URL: springerlink.com/content/9vdp5b40ep2pjvba/>, pp. 1-5 as printed.
- Haniotakis et al., "Security Enhancement Through Multiple Path Transmission in Ad Hoc Networks," IEEE Intl. Conference on Communications, Jun. 20-24, 2004, 5 pgs.
- Home et al., Escrow services and incentives in Peer-to-Peer Networks, 2001, Retrieved from the Internet URL: dl.acm.org/citation.cfm?id=501168, pp. 1-10 as printed.
- "How to burn the ISO CD image to a blank CD in Windows," XP55025538, Retrieved from the Internet: URL: <http://web.archive.org/web/20100817210432/http://webconverger.org/iso/> [retrieved on Apr. 24, 2012].
- Kubiataowicz et al., OceanStore: an architecture for global-scale persistent storage, Retrieved from the Internet <URL: dl.acm.org/citation.cfm?id=356989.357007>, pp. 1-12 as printed.
- Lee et al., "Efficient Public Key Broadcast Encryption Using Identifier of Receivers," Information Security Practice and Experience Lecture Notes in Computer Science, Springer, Berlin, DE, pp. 153-164, Jan. 1, 2006.
- Menezes et al., "Handbook of Applied Cryptography, Key Management Techniques," pp. 543-590, XP002246921, section 13.5.2(i) (Jan. 1, 1996).
- Menezes et al., "Handbook of Applied Cryptography, Key Establishment Protocols," pp. 489-541, XP002304953, section 12.7 (Jan. 1, 1997).
- Ogata et al., "Fault Tolerant Anonymous Channel," Info. and Communications Security First Intl. Conference Proceedings, ICIS '97, Beijing China; vol. 1334/1997, pp. 440-444, 1997.
- Rivest, "All-Or-Nothing Encryption and The Package Transform," Proc. of the 4th Intl. Workshop on Fast Software Encryption (1997), 9 pgs.
- Schnitzer et al., "Secured storage using secureParser," Storagess'05. Proceedings of the 2005 ACM Workshop on Storage Security and Survivability, [Proceedings of the ACM Workshop on Storage Security and Survivability. Storagess], New York, NY: ACM, US, P, Nov. 1, 2005 (Jan. 1, 2005); XP002582438, ISBN: 978-1-59593-233-4; Retrieved from the Internet: URL: <http://delivery.acm.org/10.1145/1110000/1103801/p135-schnitzer.pdf?key1=1103801&key2=6709743721&co11=GUIDE&d1=GUIDE&CFID=87774082&CFTOKEN=51146011> [retrieved on May 10, 2010].
- "Webconverger 2.30 CD versions and above," Jul. 26, 2010, XP55025541, Retrieved from the Internet: URL: <http://web.archive.org/web/20100726105024/http://webconverger.org/homepage/> [retrieved on Apr. 24, 2012].
- Johnson, "MLS-Net and Secure Parser: A New Method for Securing and Segregating Network Data [Online] Jul. 8, 2007" XP002582437 5th International Symposium on Risk Management and Informatic: WMSI2007, Orlando Florida, USA.
- Klensin, J., "Simple Mail Transfer Protocol; rfs5321.txt," Simple Mail Transfer Protocol; RFC5321.TXT, Internet Engineering Task Force, IETF; Standard, Internet Society (ISOC) 4, Rue Des Falaises CH- 1205 Geneva, Switzerland, XP015060297 (Oct. 2008) (Resubmitted).
- Koç, Erol, "Access Control in Peer-to-Peer Storage Systems," Oct. 1, 2006, XP55036196, Retrieved from the Internet: URL: <ftp://ftp.tik.ee.ethz.ch/pub/studfents/2006-So/MA-2006-10.pdf> [retrieved on Aug. 23, 2012] p. 61-62 (Resubmitted).
- Xu, "Locality Driven Key Management Architecture for Mobile Ad-hoc Networks, 2004 IEEE International Conference on Mobile Ad-hoc and Sensor Systems," pp. 436-446.

* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

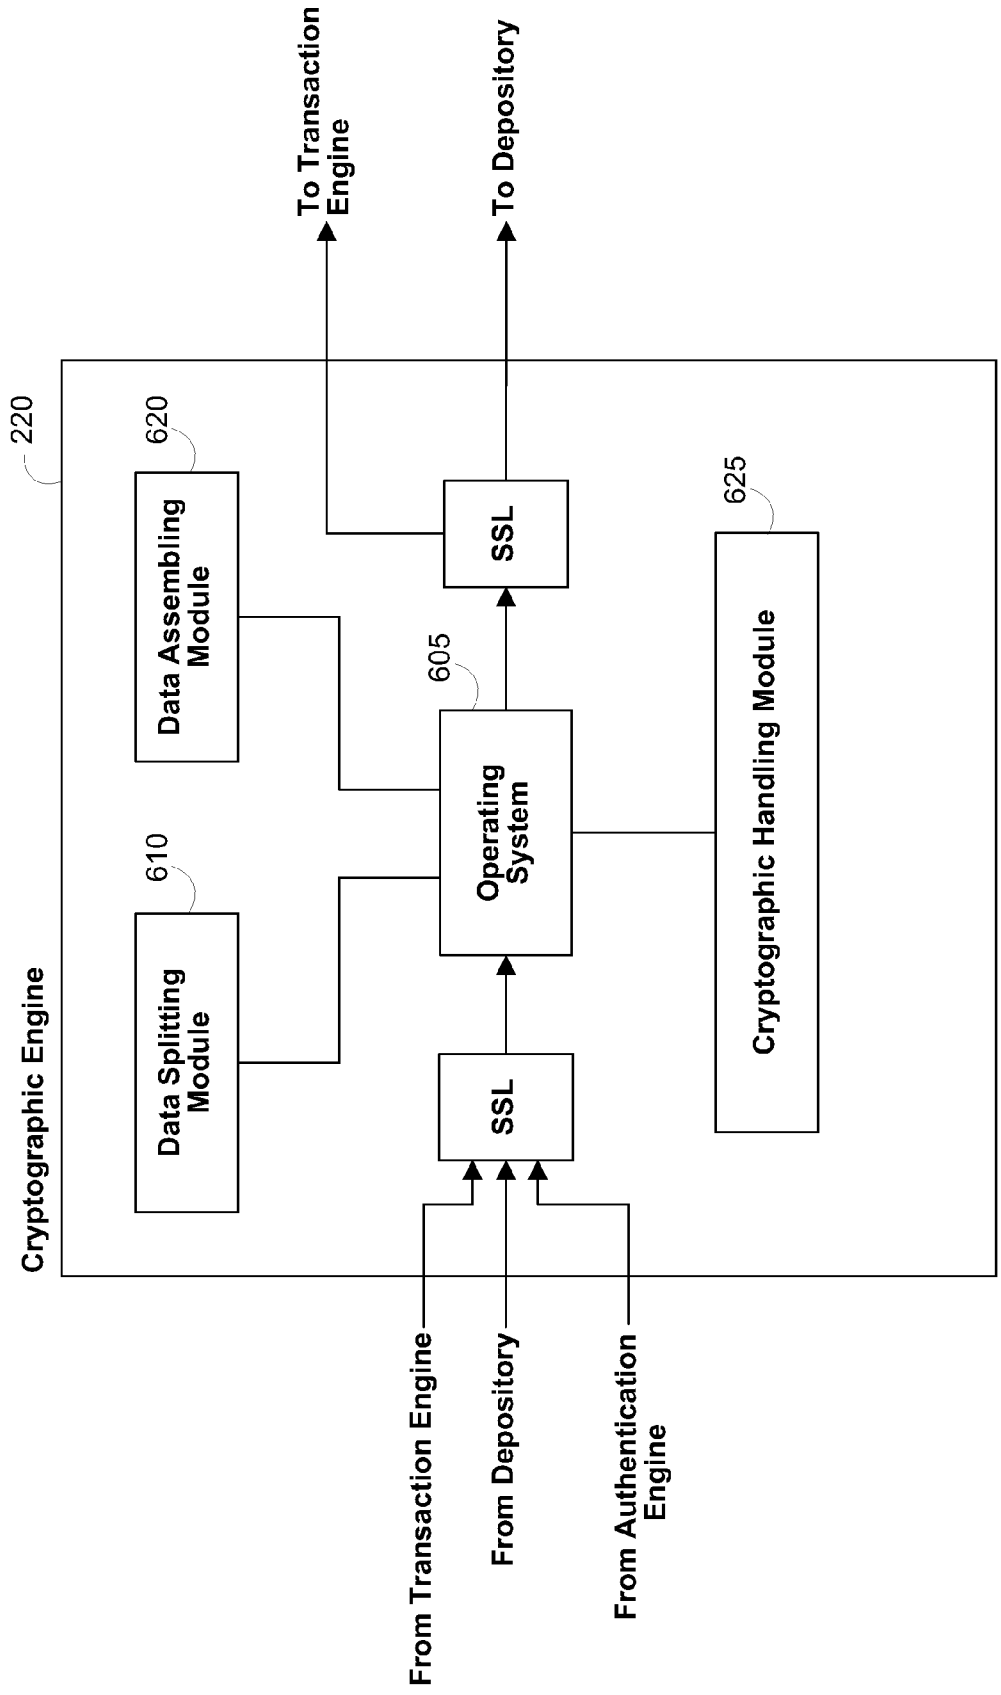


FIG. 6

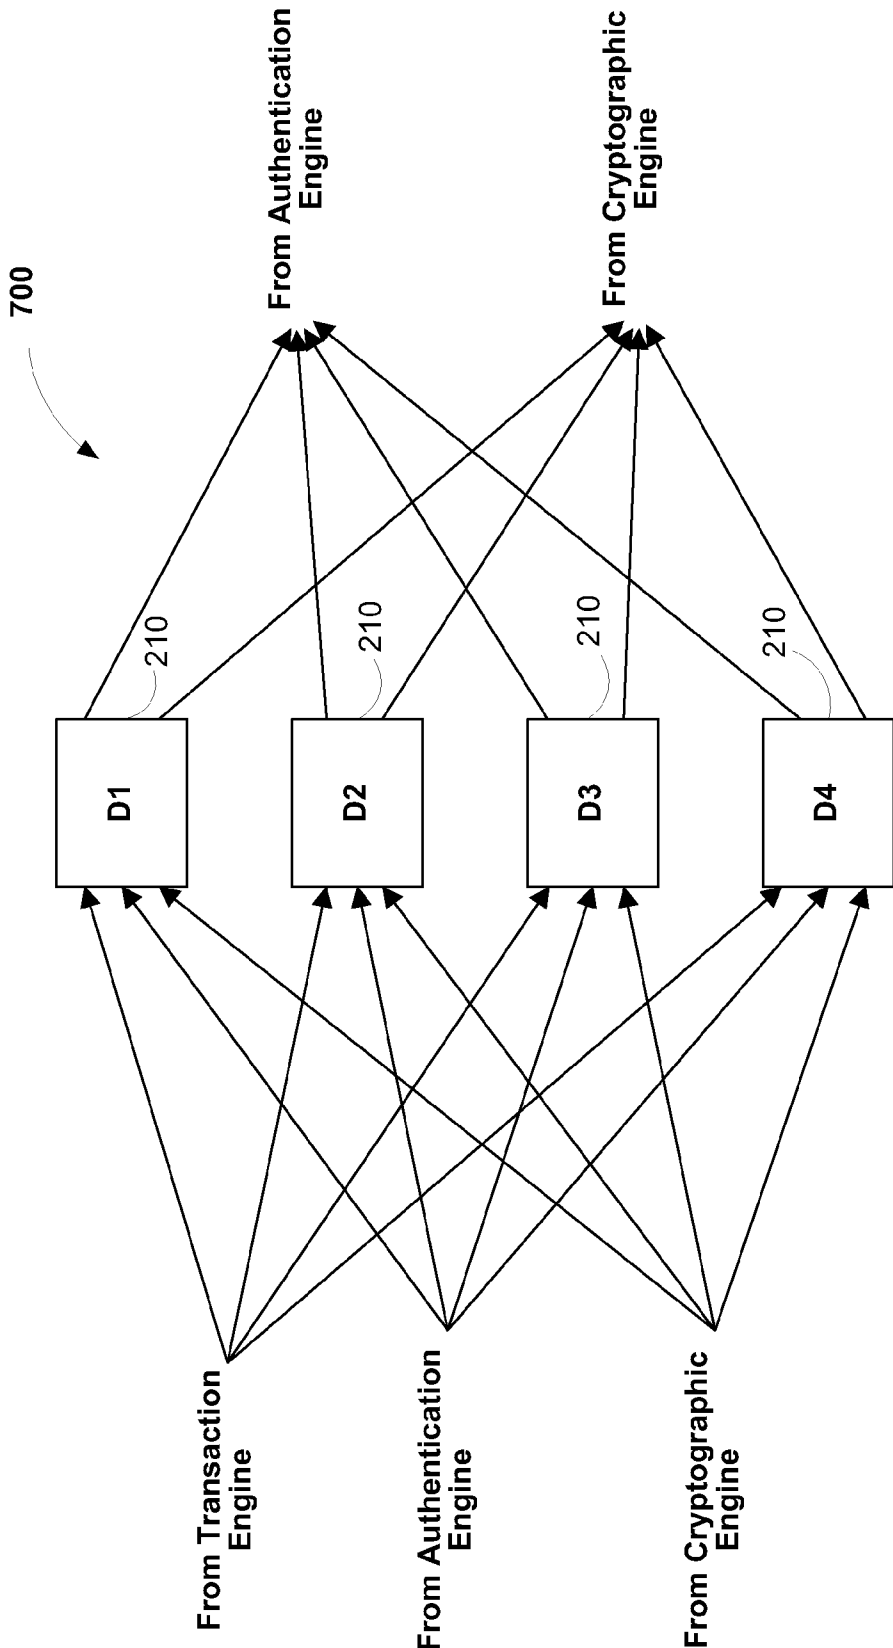


FIG. 7

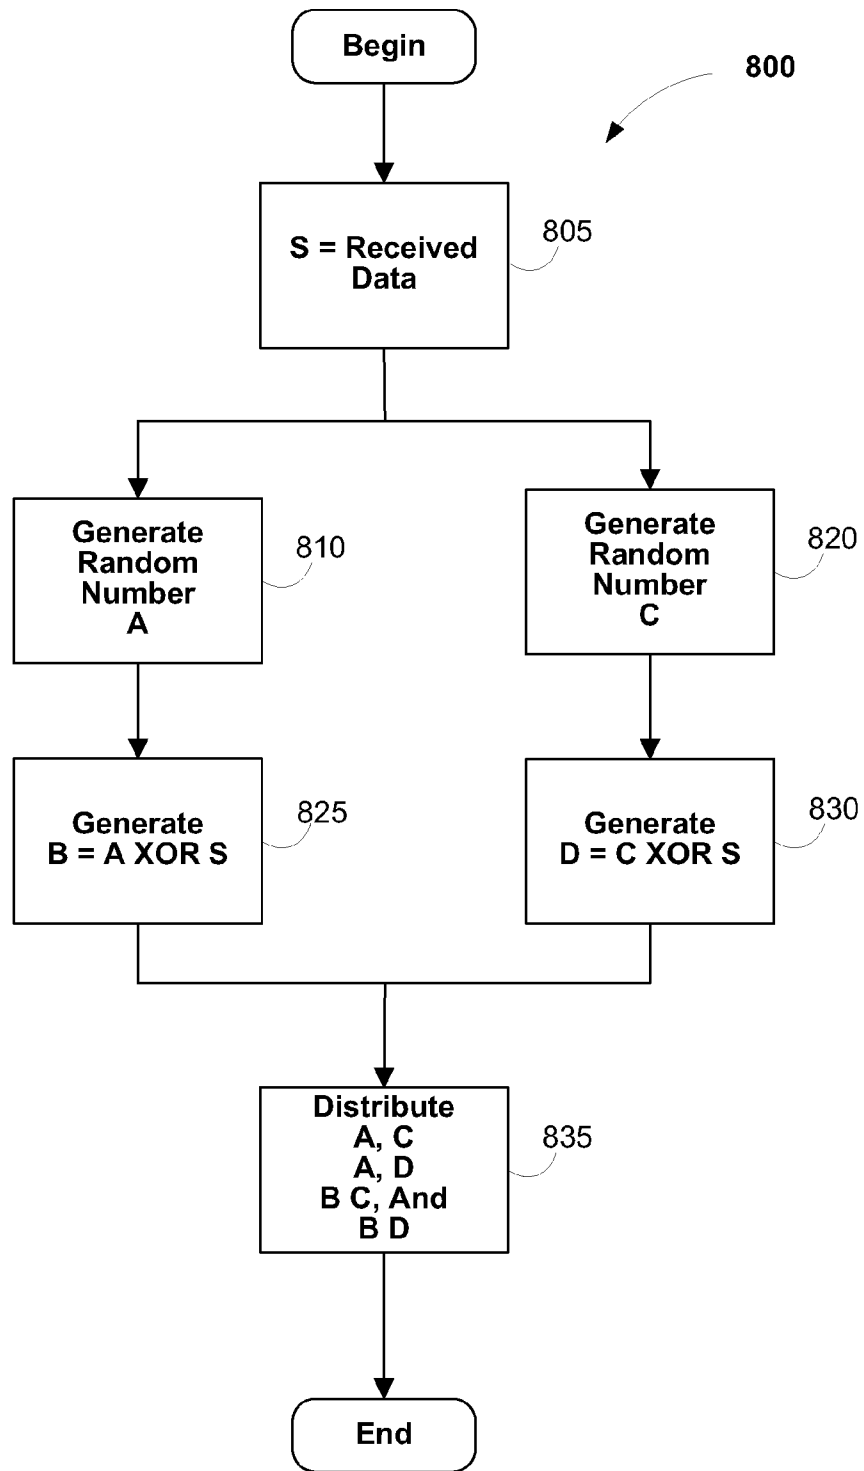



FIG. 8

900



Enrollment Data Flow			
Send	Receive	SSL	Action
User	Transaction Engine (TE)	1/2	Transmit Enrollment Authentication Data (B) and the User ID (UID) encrypted with the Public Key of the Authentication Engine (AE) as (PUB_AE(UID,B))
TE	AE	Full	Forward Transmission
			AE Decrypts and Splits Forwarded Data
AE	The Xth Depository (DX)	Full	Store Respective Portion of Data
When Digital Certificate Requested			
AE	Cryptographic Engine (CE)	Full	Request Key Generation
			CE Generates and Splits Key
CE	TE	Full	Transmit Request for Digital Certificate
TE	Certification Authority (CA)	1/2	Transmit Request
CA	TE	1/2	Transmit Digital Certificate
TE	User	1/2	Transmit Digital Certificate
TE	MS	Full	Store Digital Certificate
CE	DX	Full	Store Respective Portion of Key

FIG. 9, Panel A

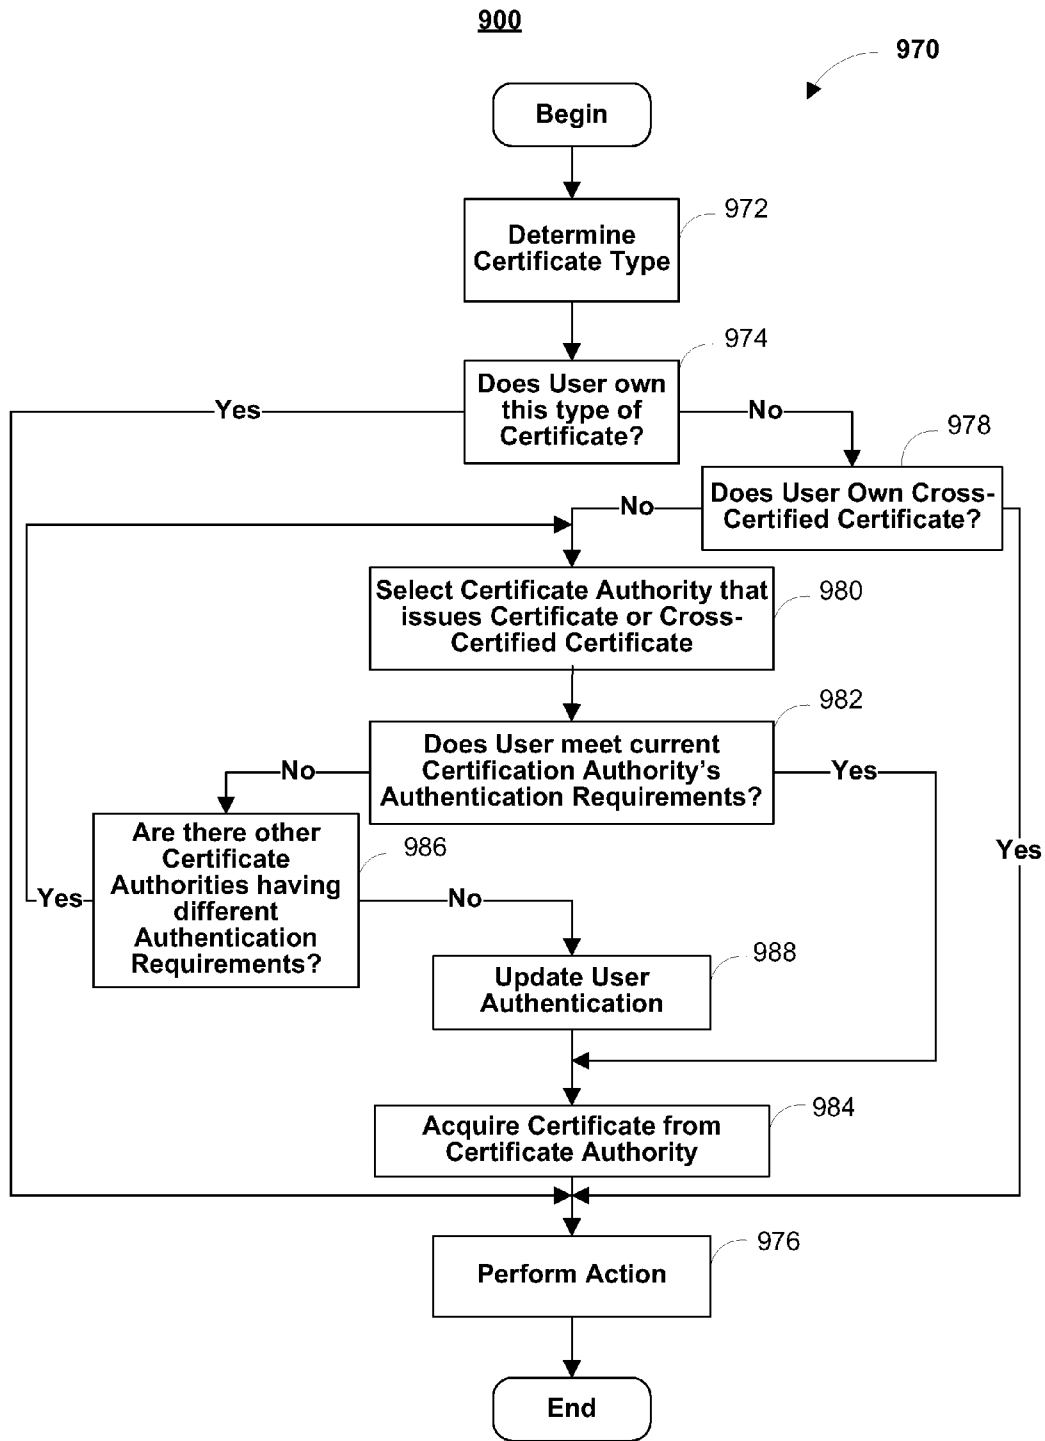


FIG. 9, Panel B

1000
▼

Authentication Data Flow				
	SEND	RECEIVE	SSL	ACTION
1005	User	Vendor	1/2	Transaction occurs, such as selecting purchase
1010	Vendor	User	1/2	Transmit transaction ID (TID) and authentication request (AR)
				Authentication data (B') is gathered from User
1015	User	TE	1/2	Transmit TID and B' wrapped in the Public Key of the Authentication Engine (AE), as (PUB_AE(TID, B'))
1020	TE	AE	Full	Forward transmission
				Enrollment authentication data (B) is requested and gathered
1025	Vendor	Transaction Engine (TE)	Full	Transmits TID, AR
1030	TE	Mass Storage (MS)	Full	Create Record in database
1035	TE	The Xth Depository (DX)	Full	UID, TID
1040	DX	AE	Full	Transmit the TID and the portion of the authentication data stored at enrollment (BX) as (PUB_AE(TID, BX))
1045				AE assembles B and compares to B'
1050	AE	TE	Full	TID, the filled in AR
1055	TE	Vendor	Full	TID, Yes/No
	TE	User	1/2	TID, confirmation message

FIG. 10

1100

Signing Data Flow				
SEND	RECEIVE	SSL	ACTION	
User	Vendor	1/2	Transaction occurs, such as agreeing on a deal	
Vendor	User	1/2	Transmit transaction identification number (TID), authentication request (AR), and agreement or message (M)	
			Current authentication data (B') and a hash of the message received by the User (h(M')) is gathered from User	
User	TE	1/2	Transmit TID, B', AR, and h(M') wrapped in the Public Key of the Authentication Engine (AE), as (PUB_AE(TID, B', h(M')))	
TE	AE	Full	Forward transmission	
			Gather enrollment authentication data	
Vendor	Transaction Engine (TE)	Full	Transmits UID, TID, AR, and a hash of the message (h(M')).	
TE	Mass Storage (MS)	Full	Create Record in database	
TE	The Xth Depository (DX)	Full	UID, TID	
DX	AE	Full	Transmit the TID and the portion of the authentication data stored at Enrollment (BX), as (PUB_AE(TID, BX))	
			The original vendor message is transmitted to the AE	
TE	AE	Full	Transmit h(M)	
1103	AE assembles B, compares to B' and compares h(M) to h(M')			
1105	AE	Cryptographic Engine (CE)	Full	Request for digital signature and a message to be signed, for example, the hashed message
1110	AE	DX	Full	TID, signing UID
1115	DX	CE	Full	Transmit the portion of the Cryptographic Key corresponding to the signing party
1120	CE assembles key and signs			
1125	CE	AE	Full	Transmit the digital signature (S) of signing party
1130	AE	TE	Full	TID, the filled in AR, h(M), and S
1135	TE	Vendor	Full	TID, a receipt=(TID, Yes/No, and S), and the digital signature of the trust engine, for example, a hash of the receipt encrypted with the trust engine's Private Key (Priv_TE(h(receipt)))
1140	TE	User	1/2	TID, confirmation message

FIG. 11

1200

Encryption/Decryption Data Flow			
Send	Receive	SSL	Action
Decryption			
			Perform Authentication Data Process 1000, include the Session Key (sync) in the AR, where the sync has been encrypted with the Public Key of the User as PUB_USER(SYNC)
			Authenticate the User
AE	CE	Full	Forward PUB_USER(SYNC) to CE
AE	DX	Full	UID, TID
DX	CE	Full	Transmit the TID and the portion of the Private Key as (PUB_AE(TID, KEY_USER))
			CE assembles the Cryptographic Key and decrypts the sync
CE	AE	Full	TID, the filled in AR including decrypted sync
AE	TE	Full	Forward to TE
TE	Requesting APP/Vendor	1/2	TID, Yes/No, Sync
Encryption			
Requesting APP/Vendor	TE	1/2	Request for Public Key of User
TE	MS	Full	Request Digital Certificate
MS	TE	Full	Transmit Digital Certificate
TE	Requesting APP/Vendor	1/2	Transmit Digital Certificate

FIG. 12

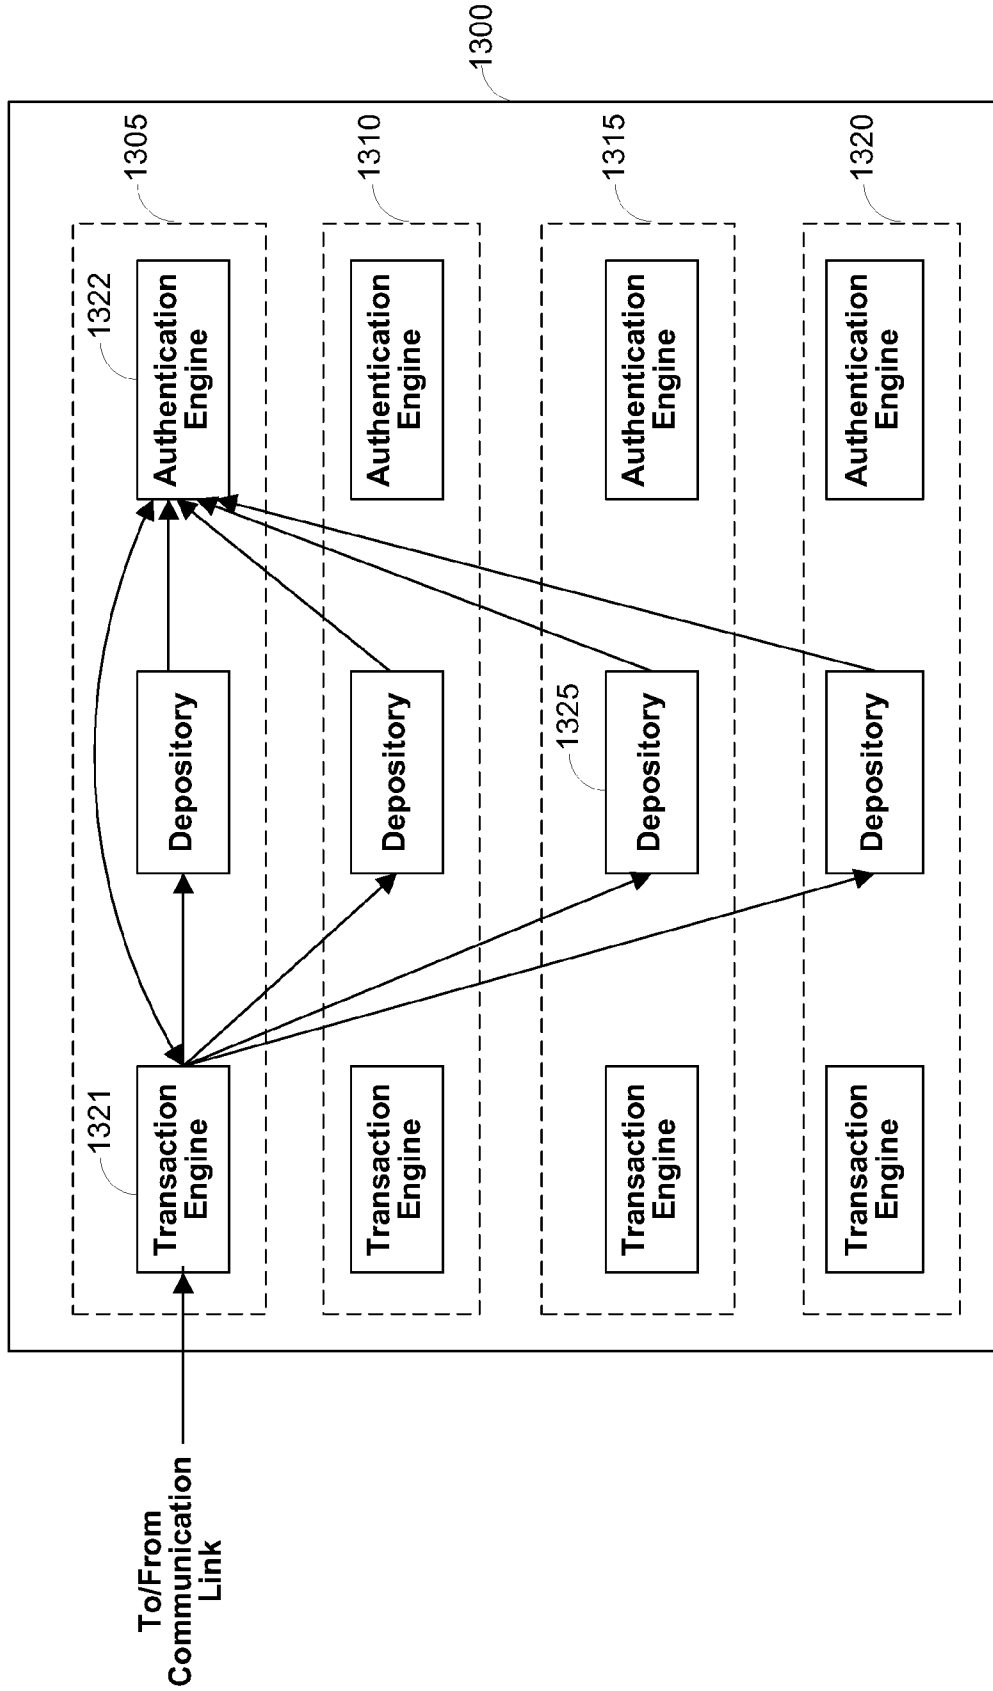


FIG. 13

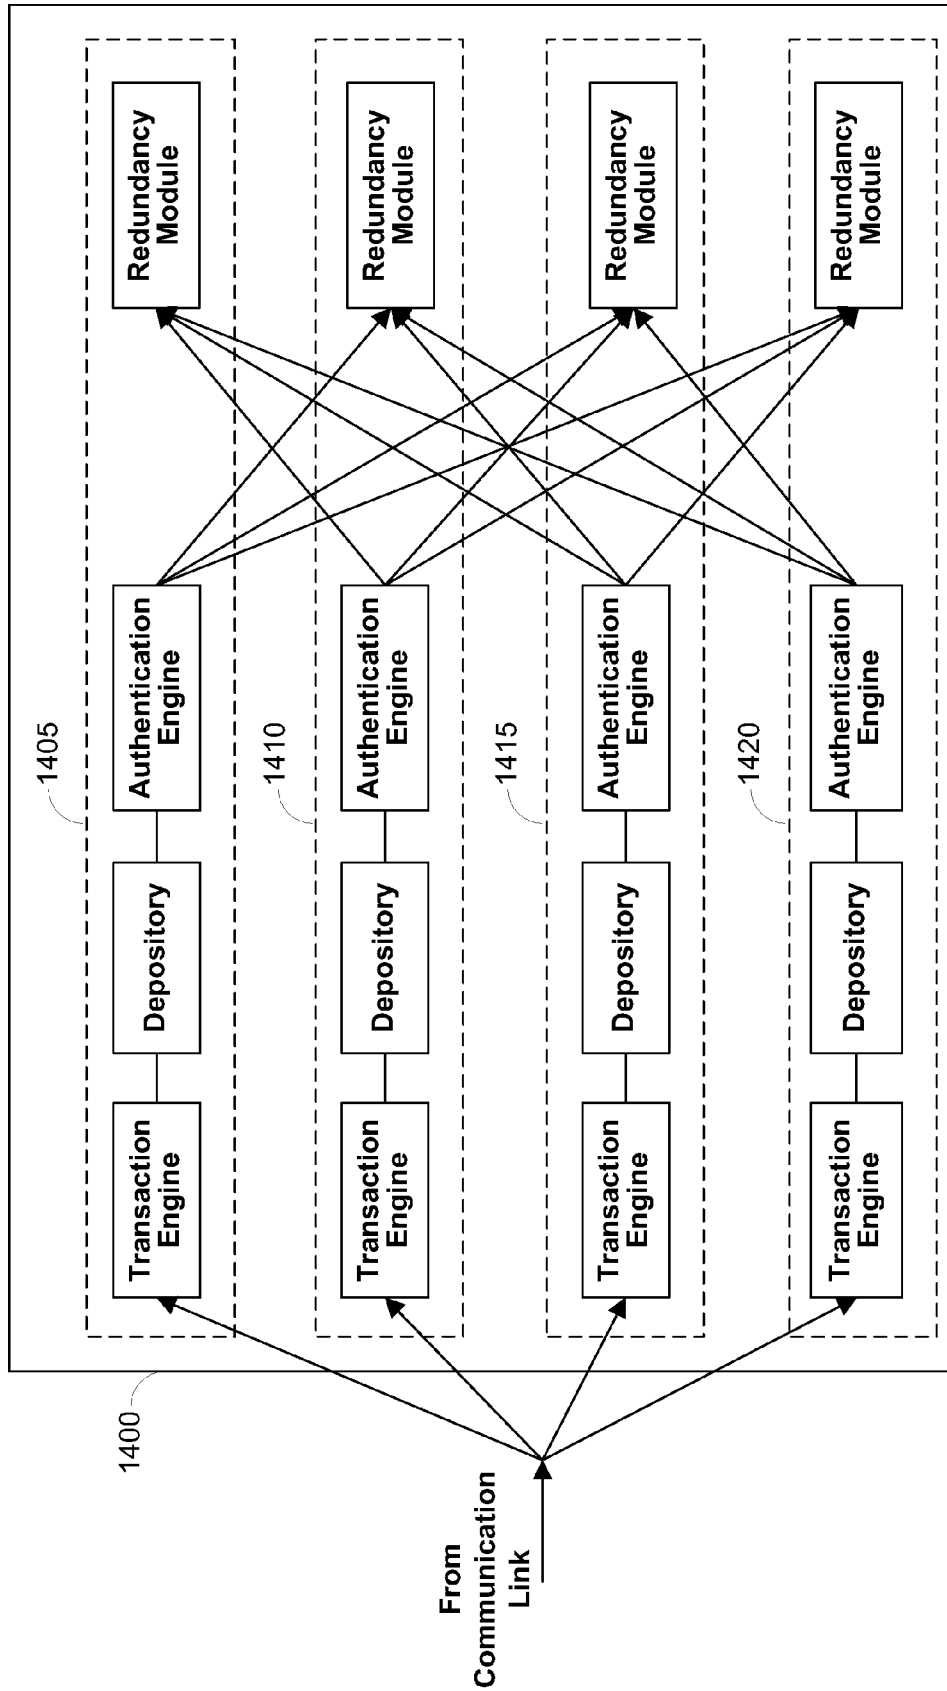


FIG. 14




FIG. 15

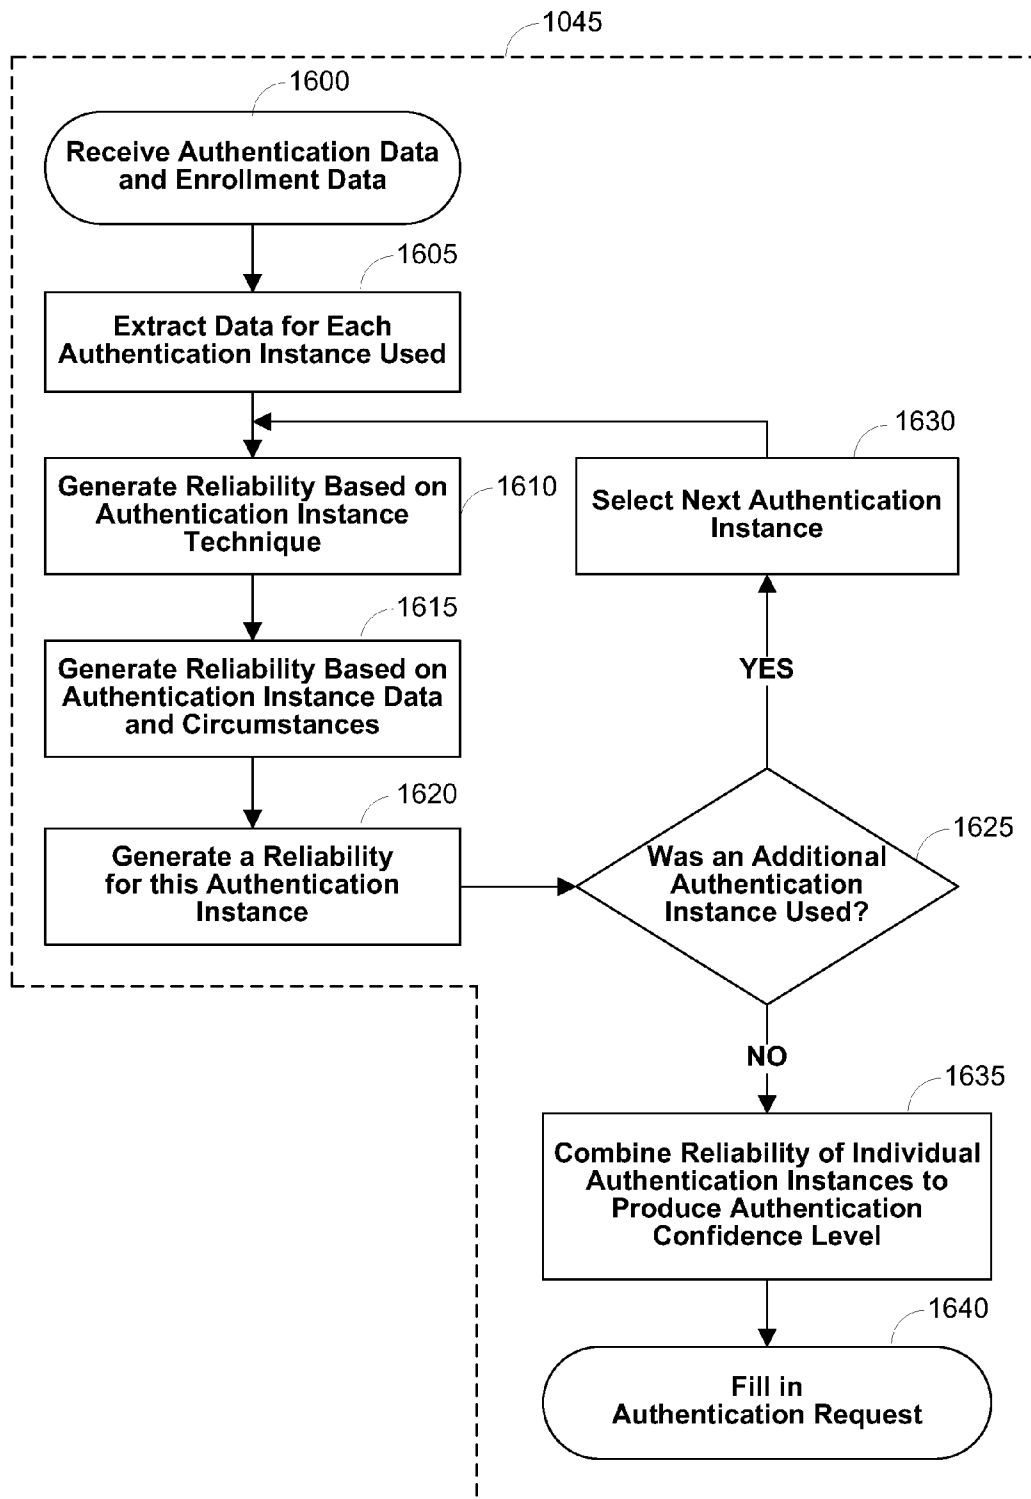


FIG. 16

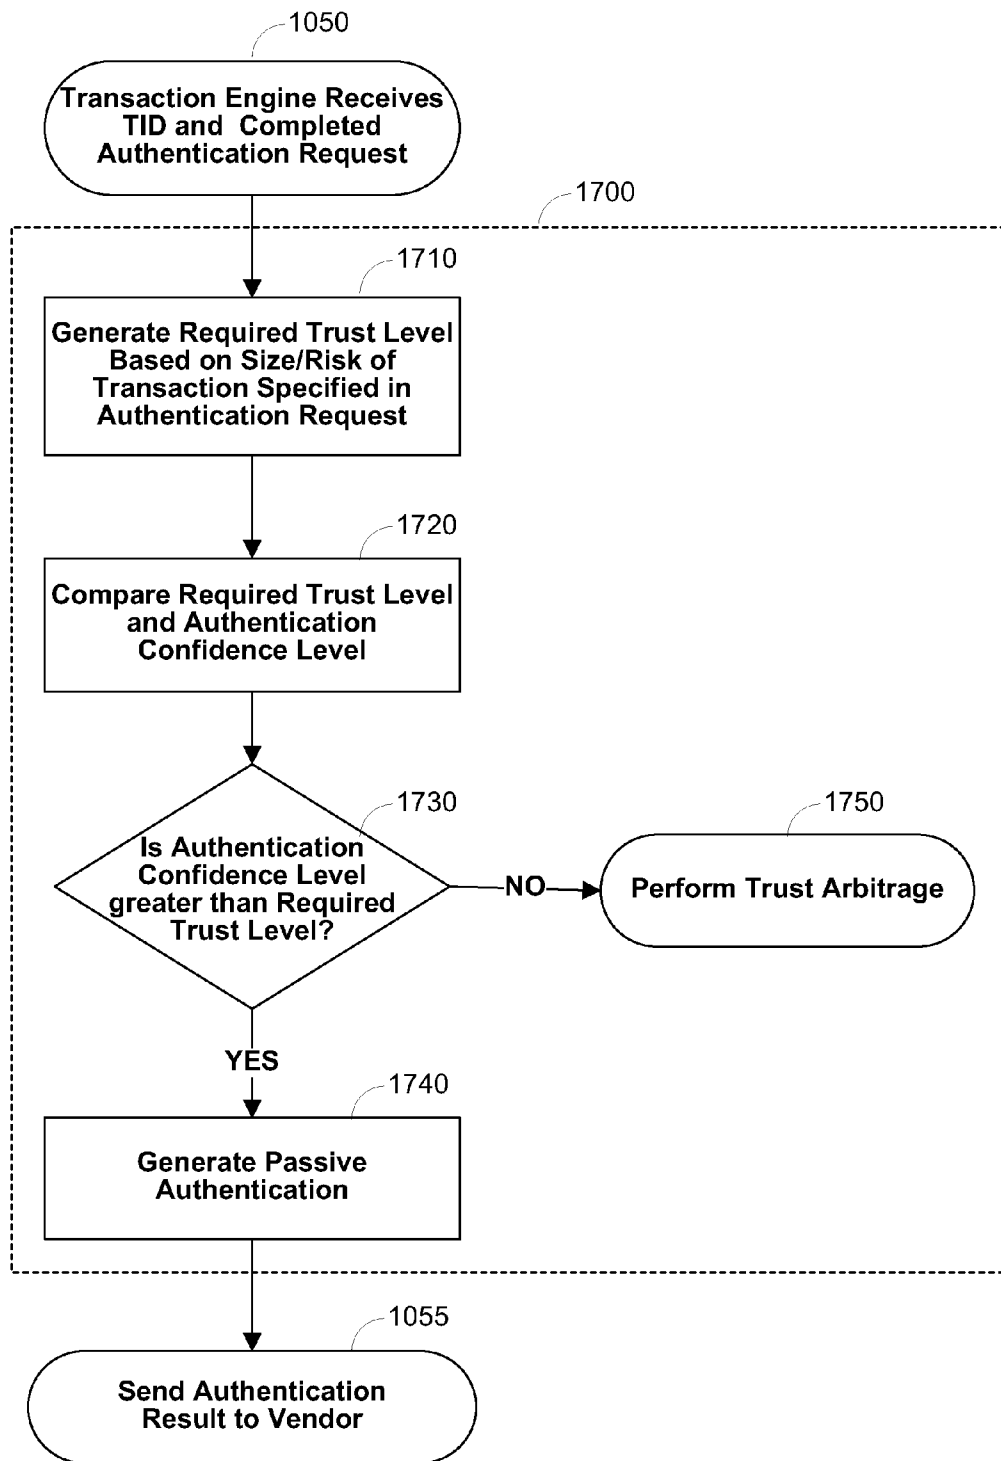


FIG. 17

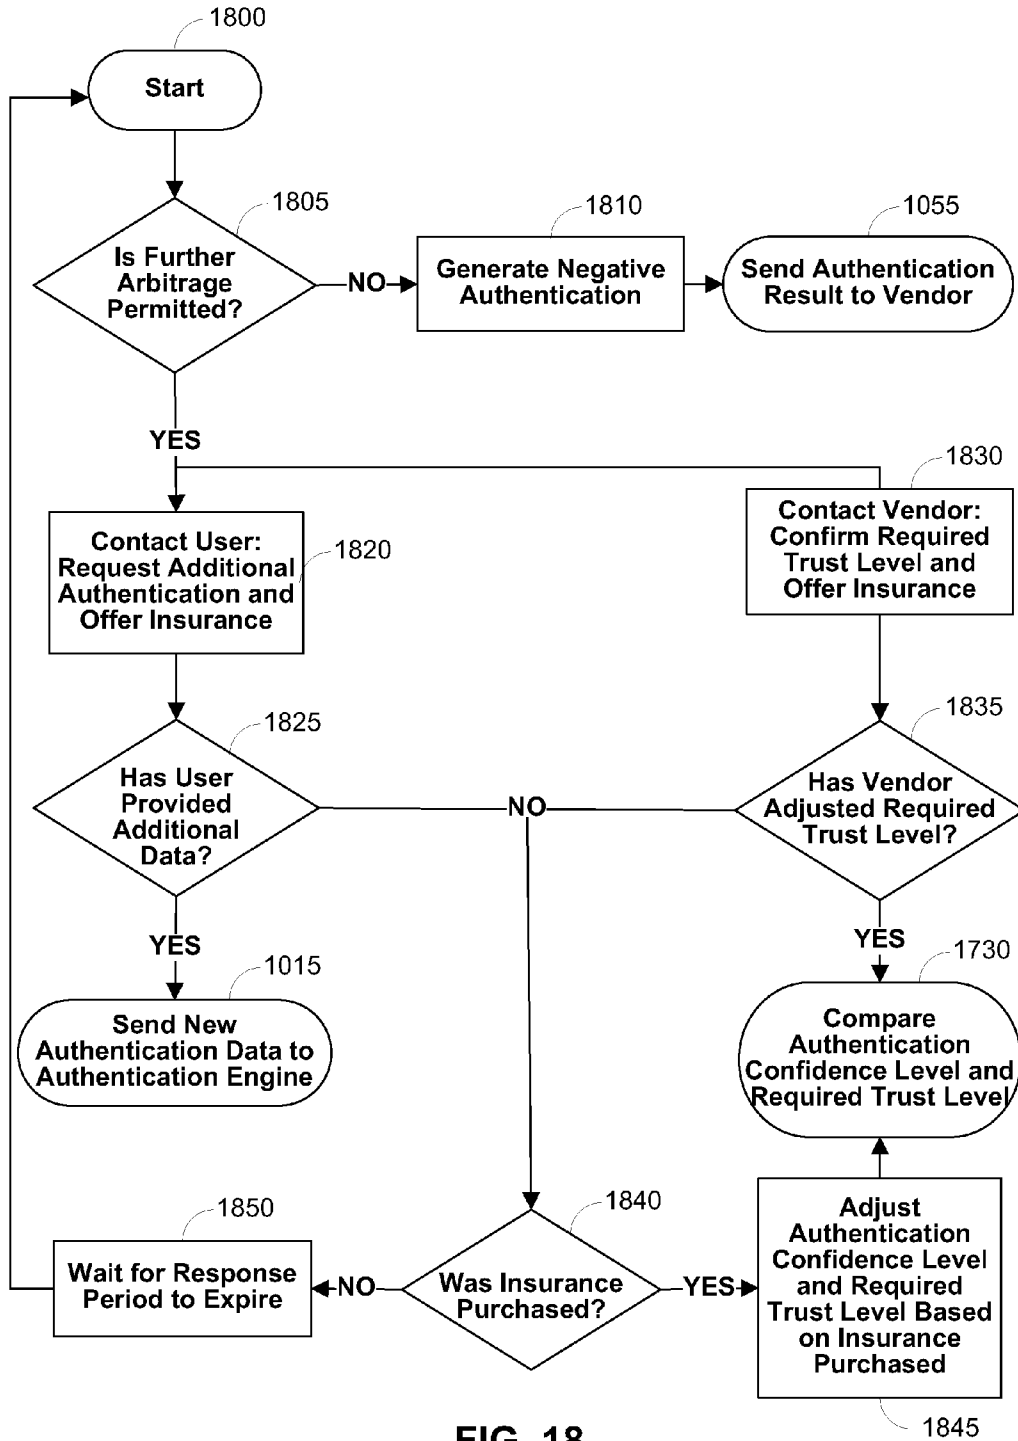


FIG. 18

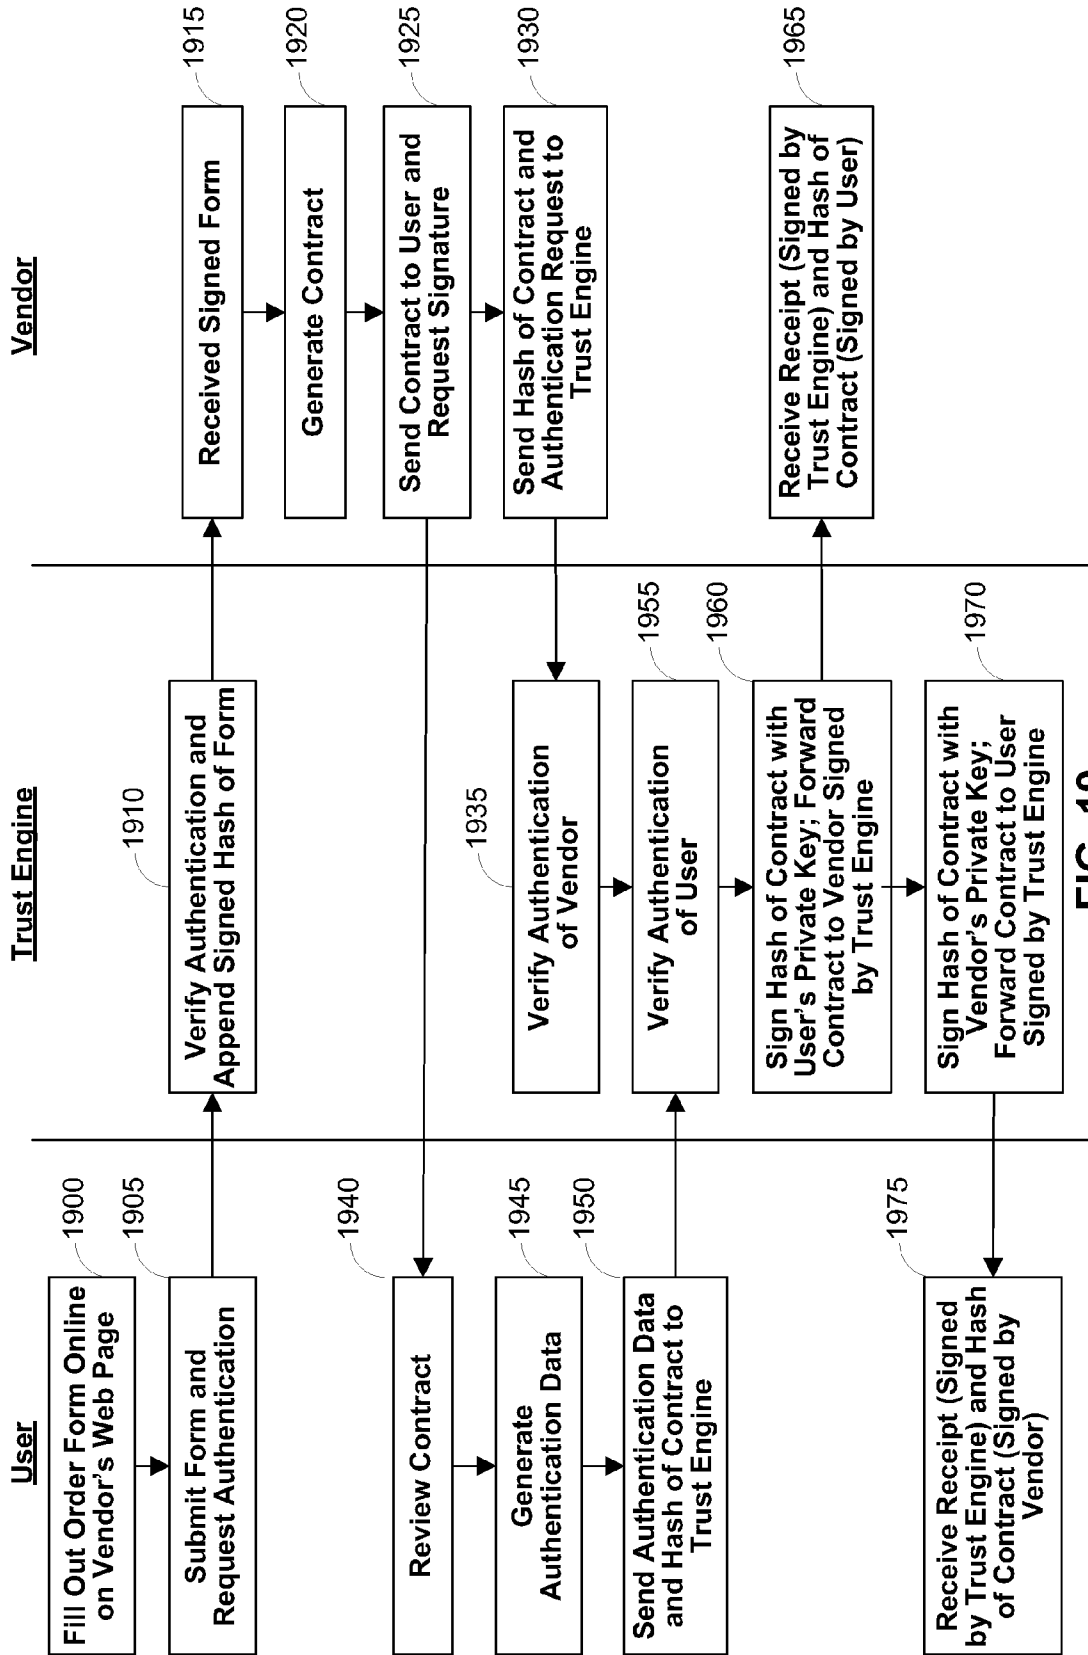


FIG. 19

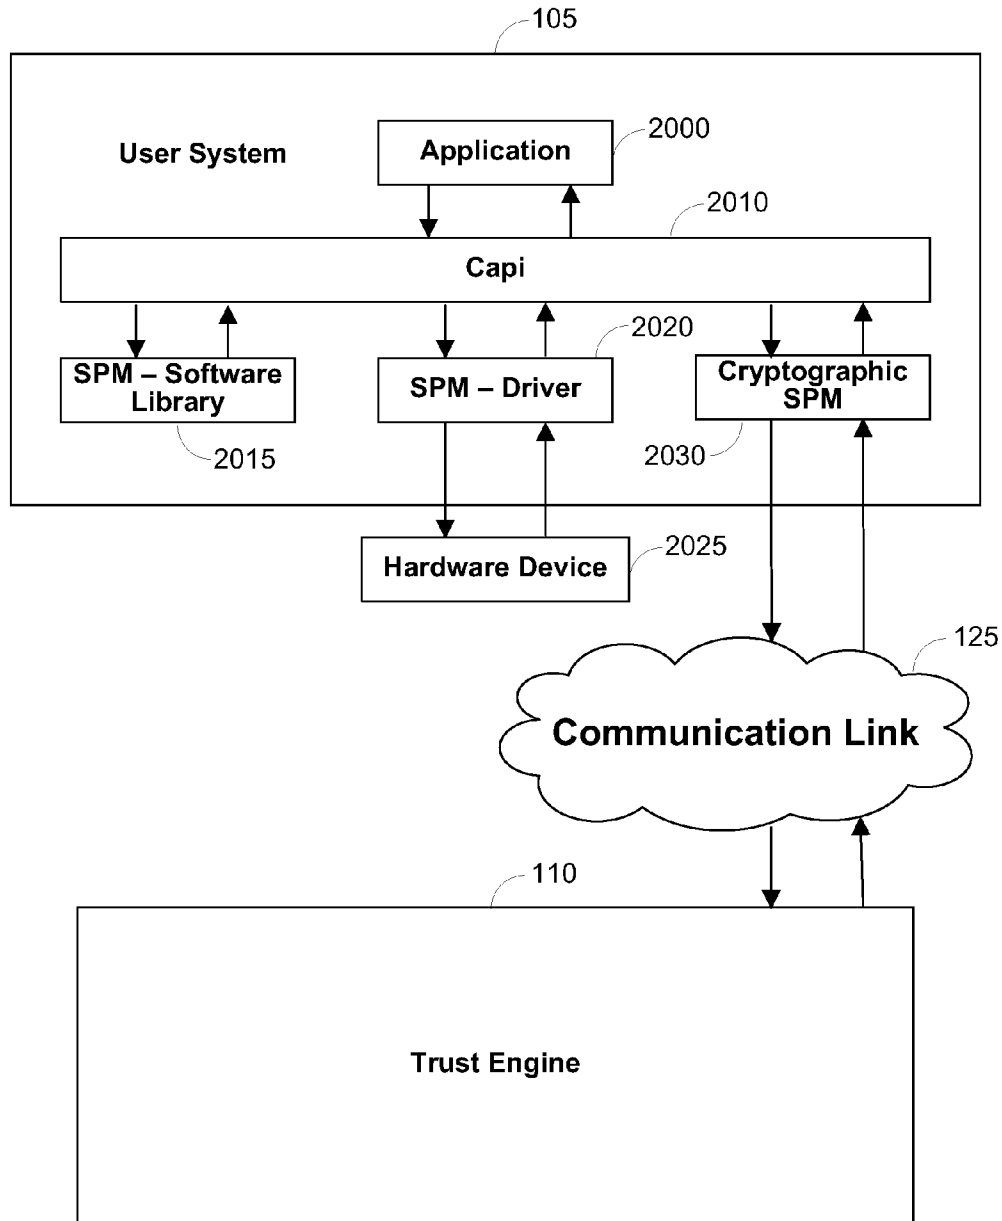


FIG. 20

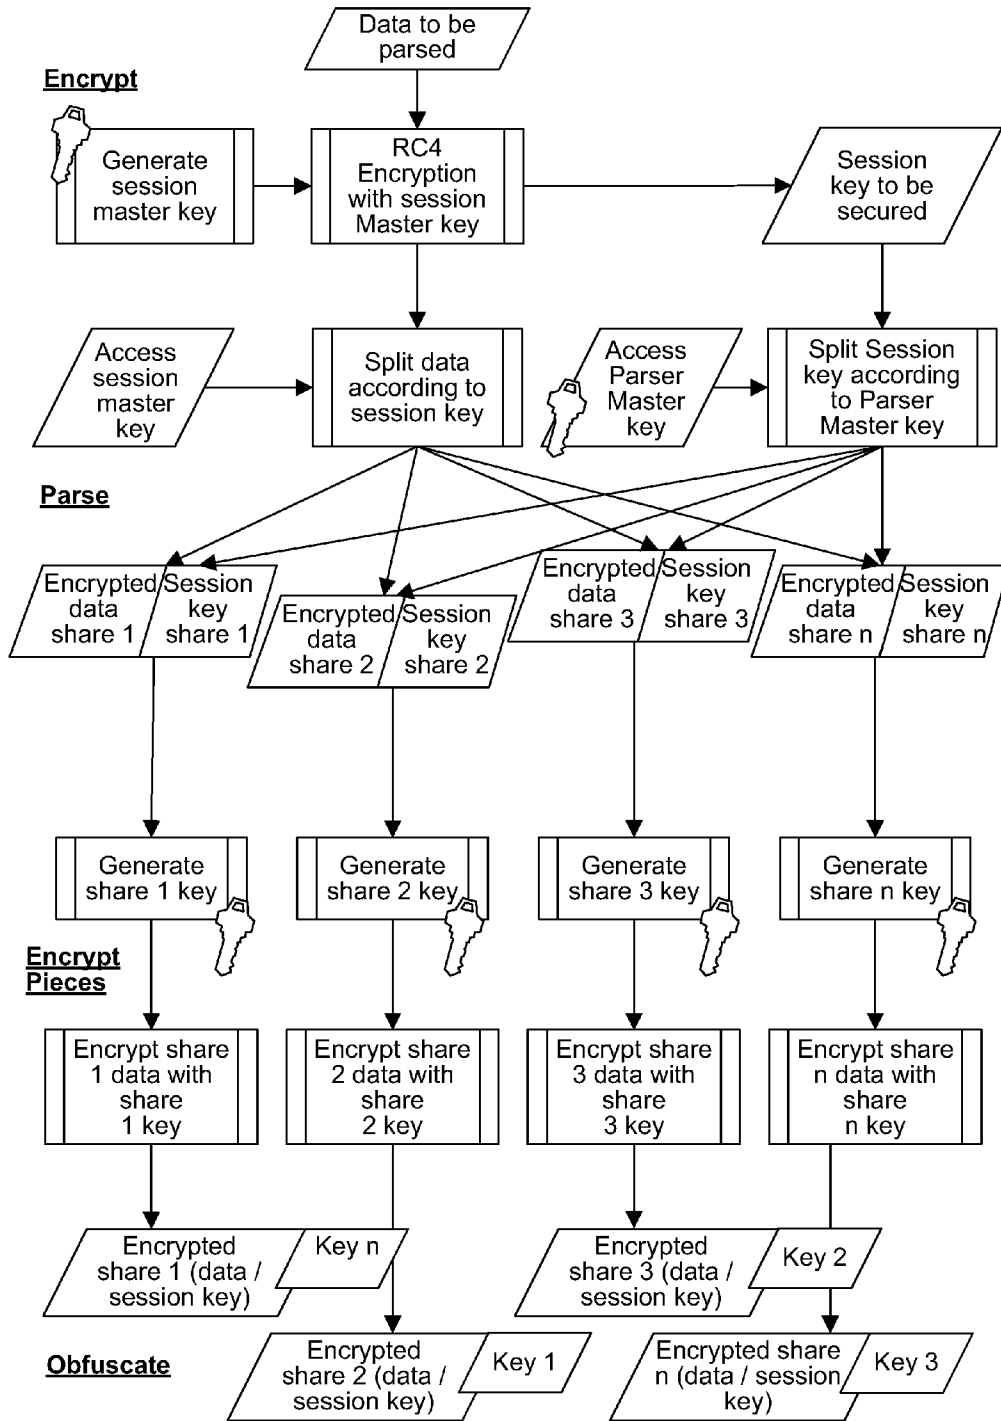


FIG. 21

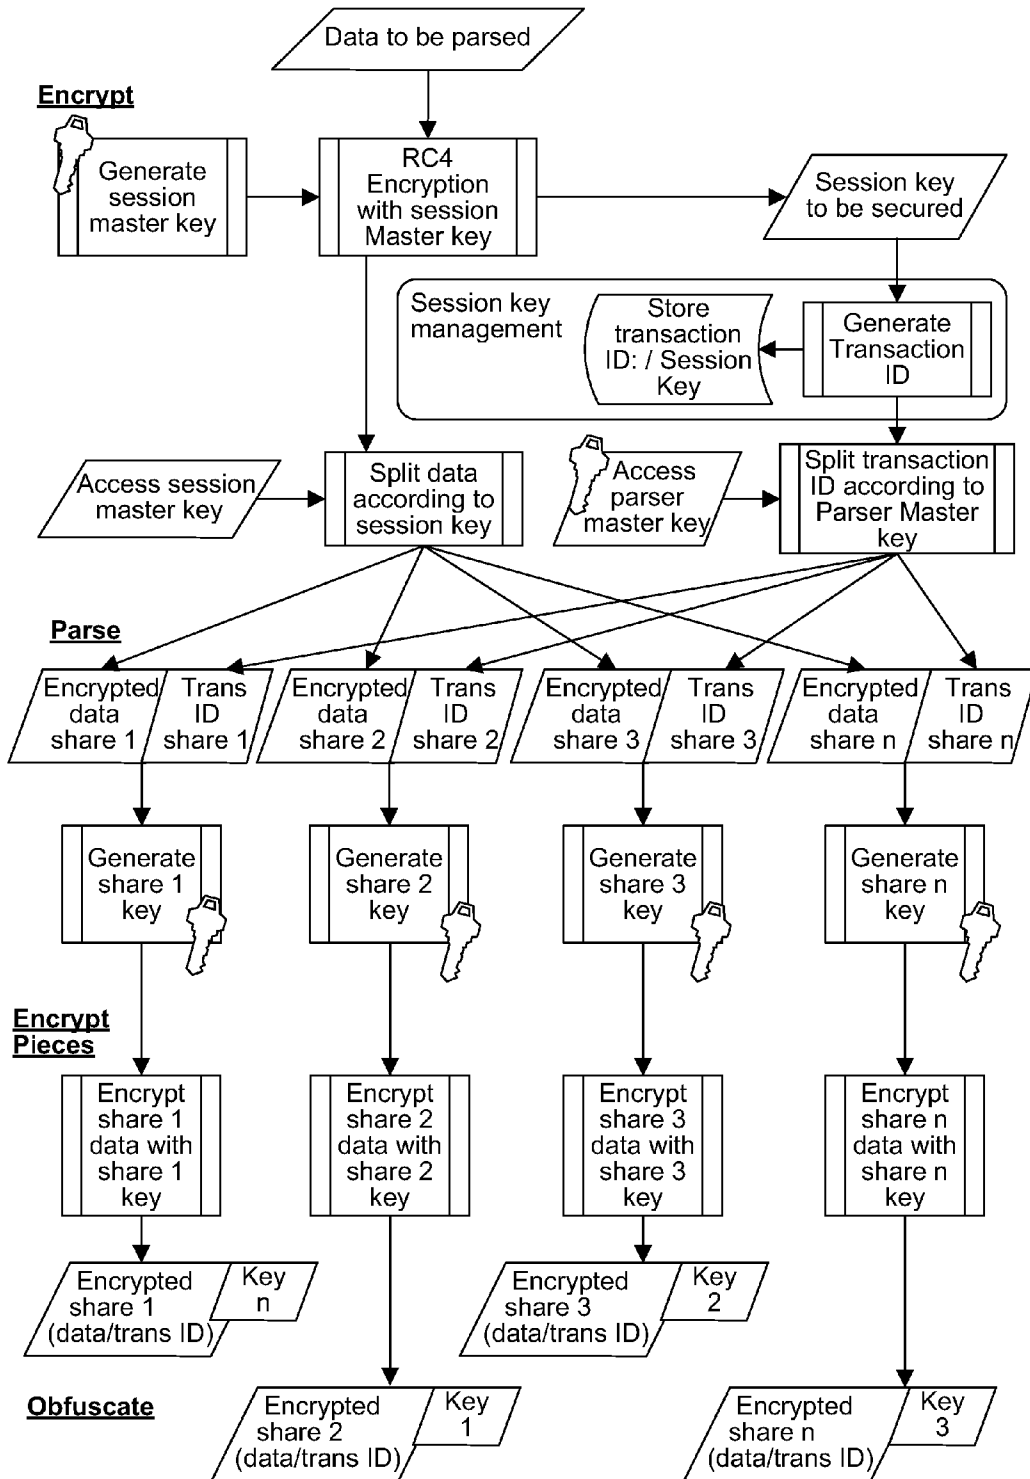


FIG. 22

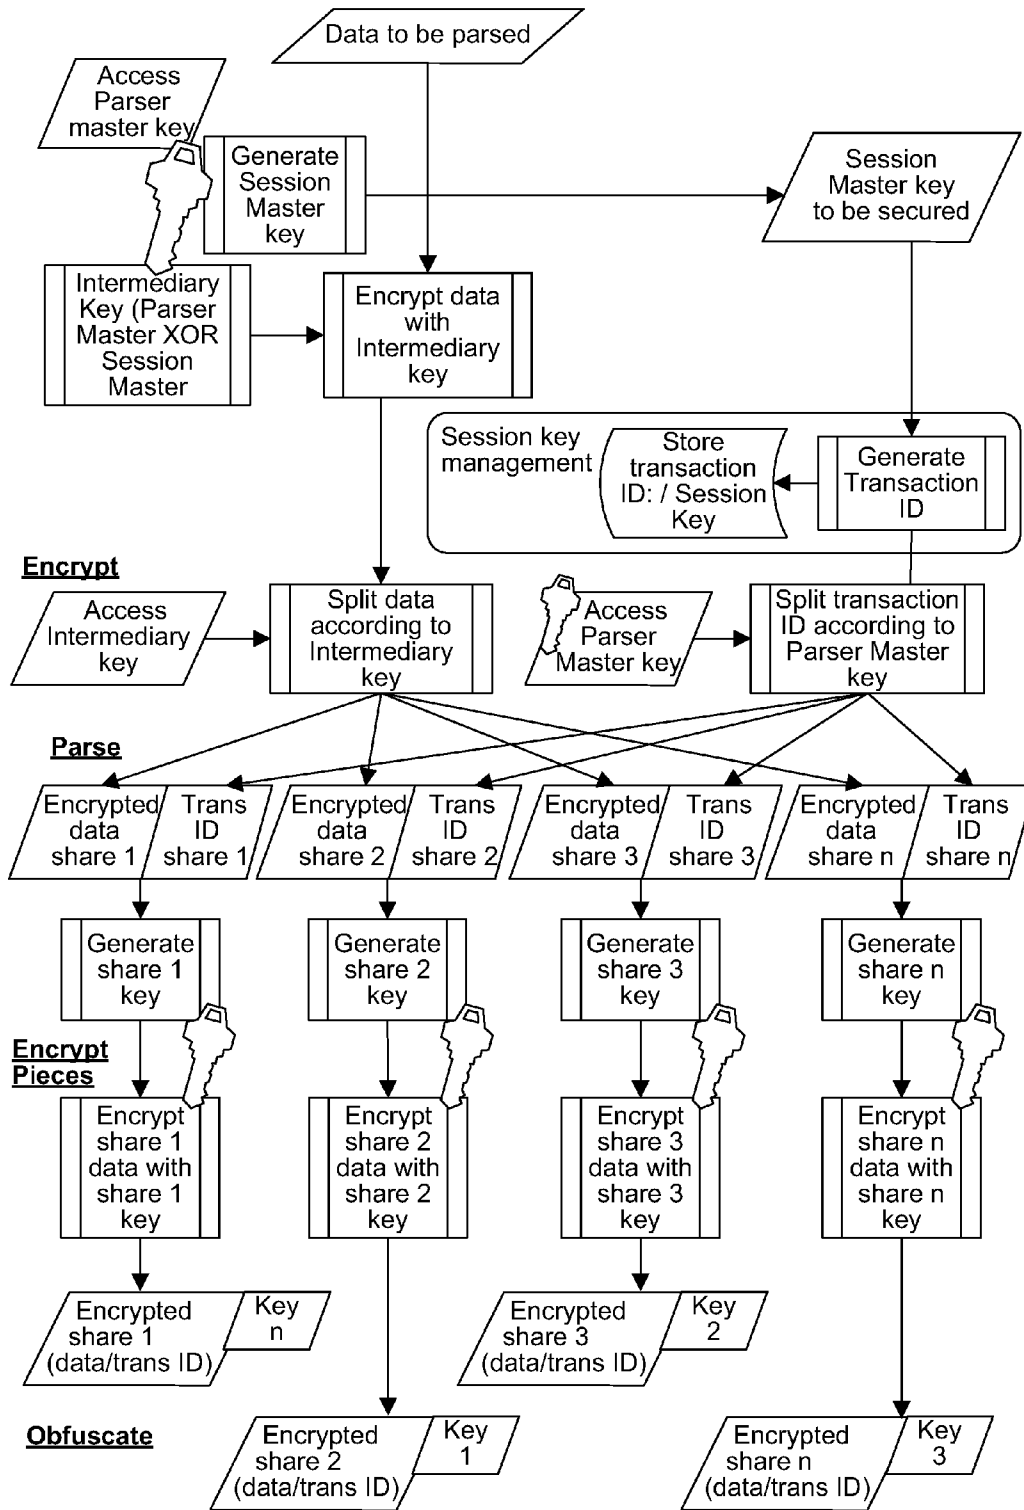


FIG. 23

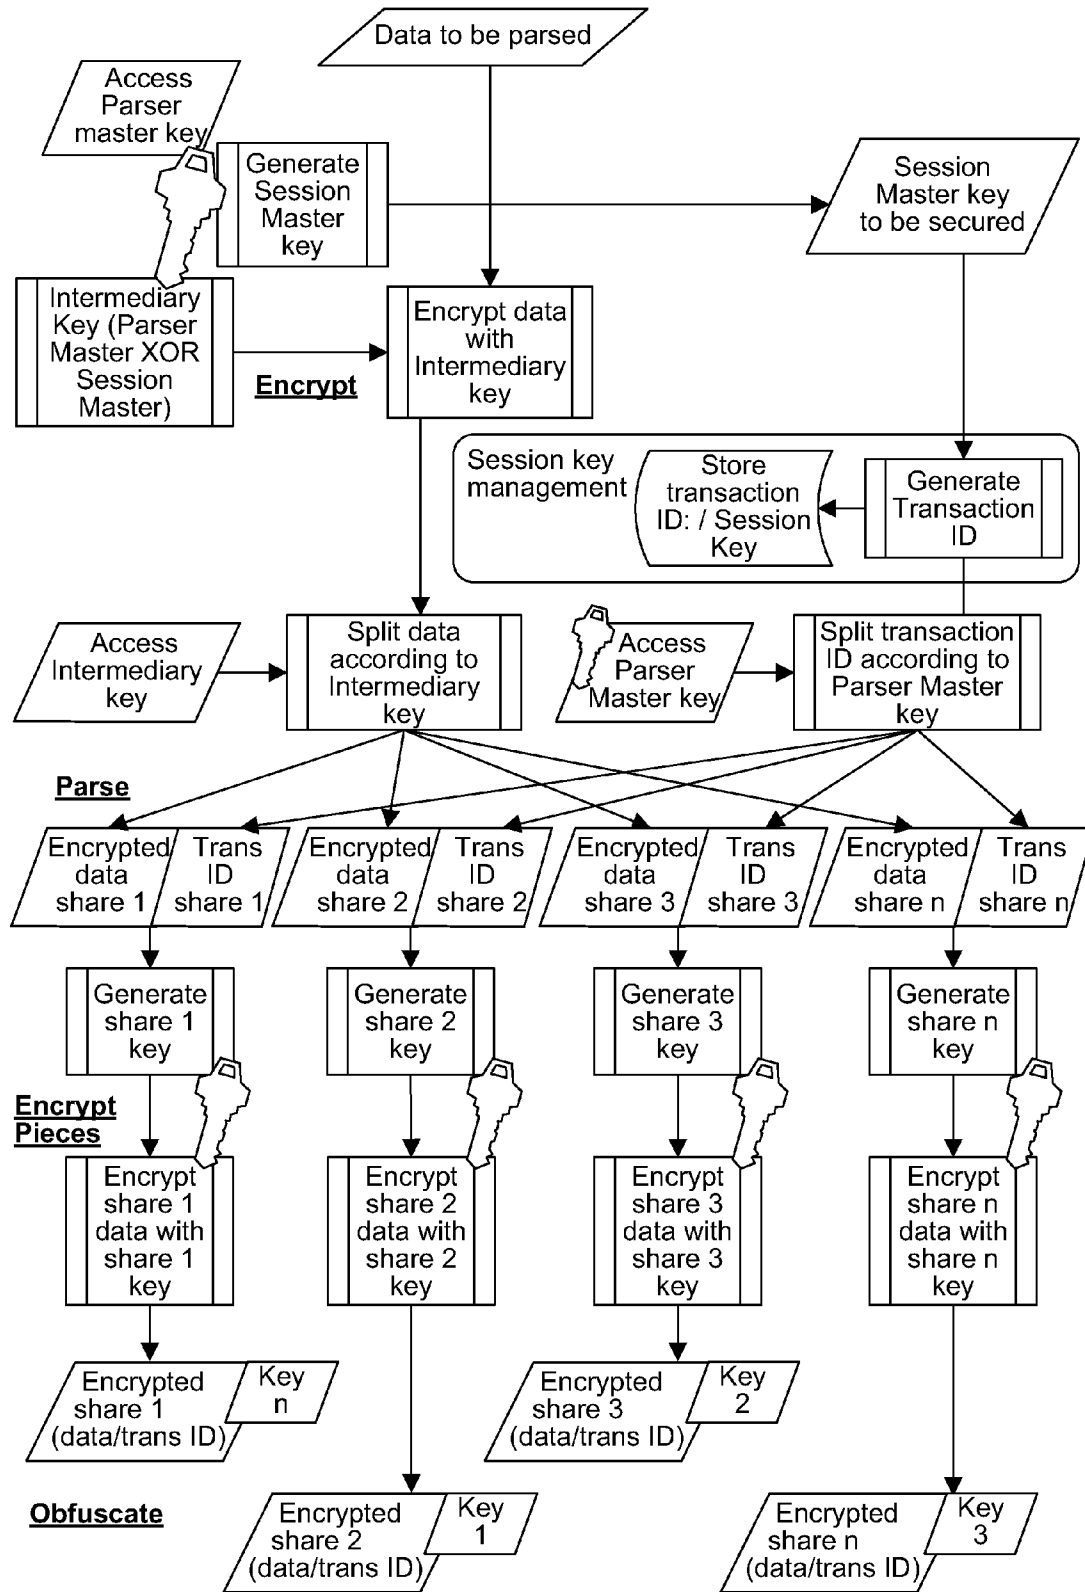


FIG. 24

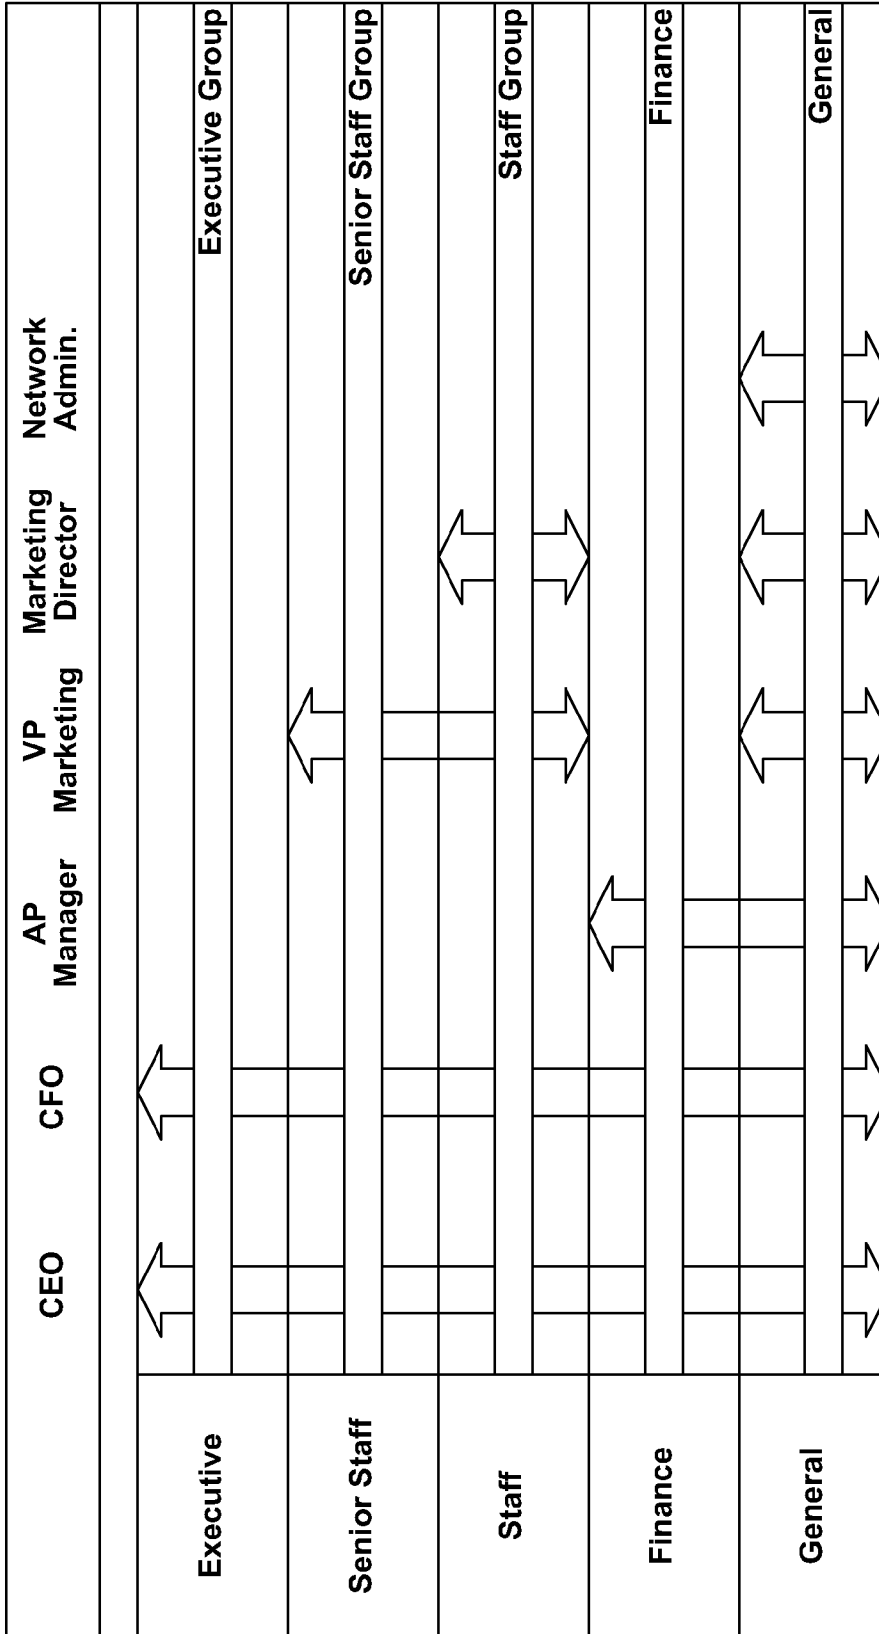


FIG. 25




FIG. 26




FIG. 27




FIG. 28




FIG. 29




FIG. 30

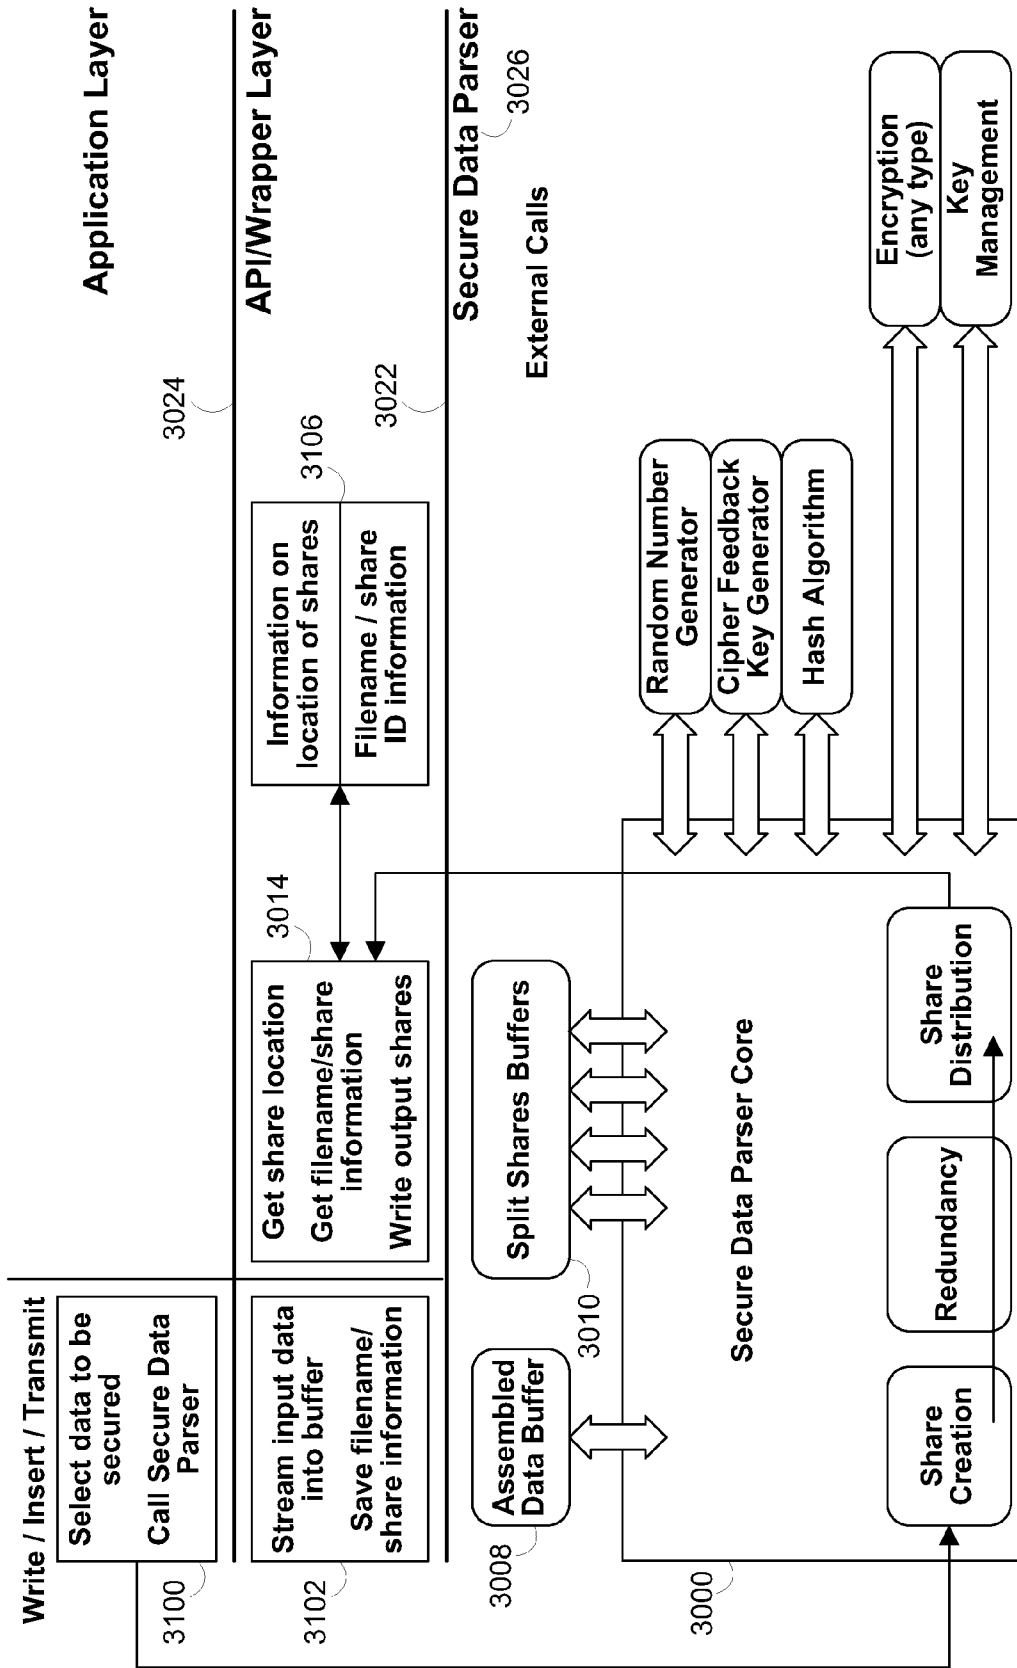


FIG. 31

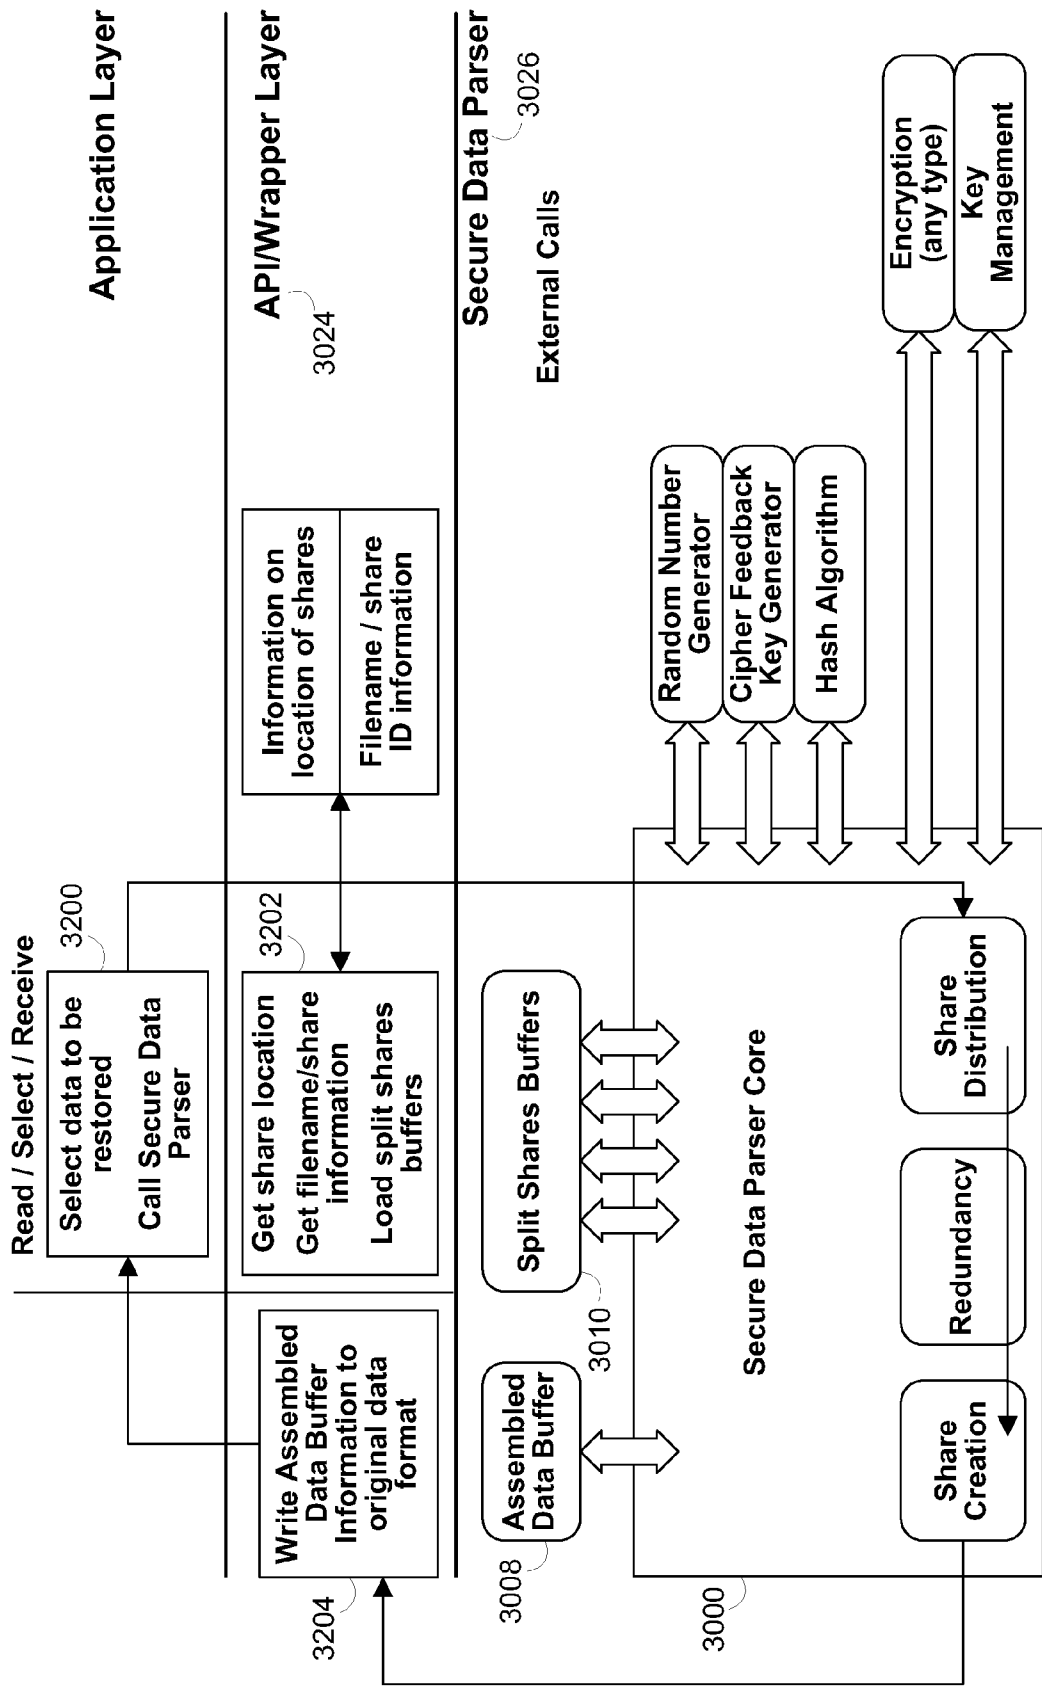


FIG. 32




FIG. 33




FIG. 34

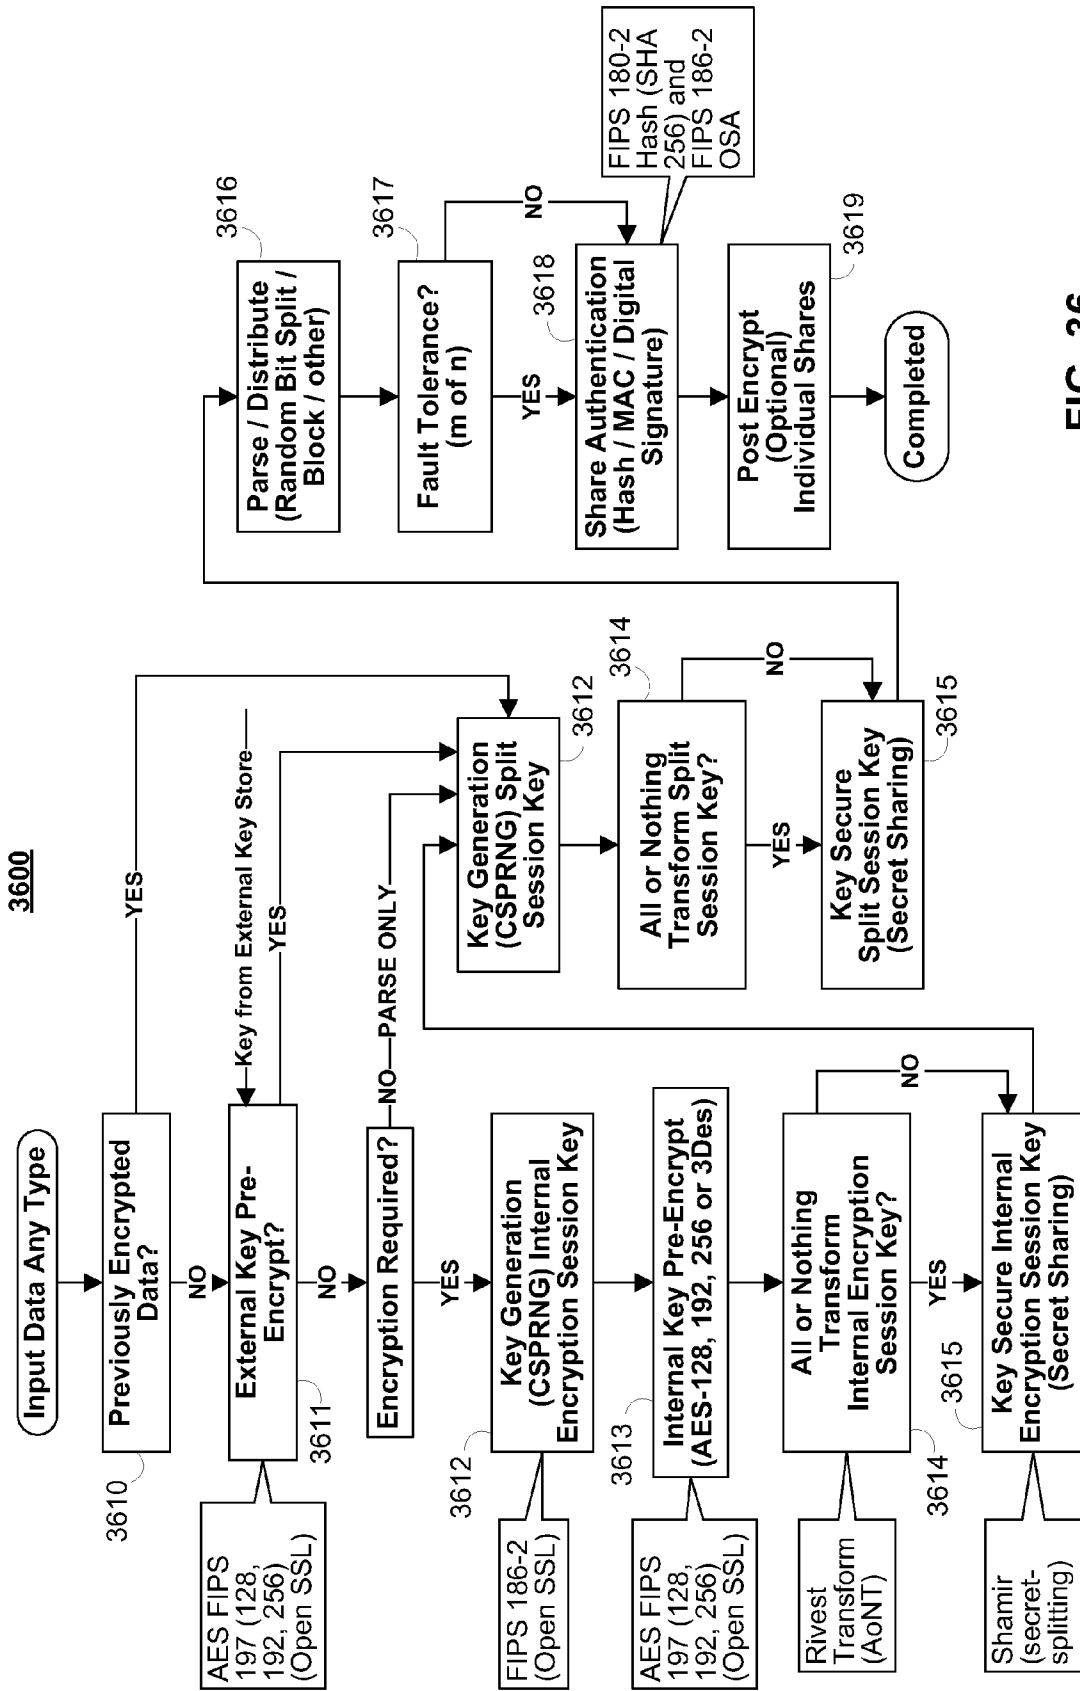


FIG. 36

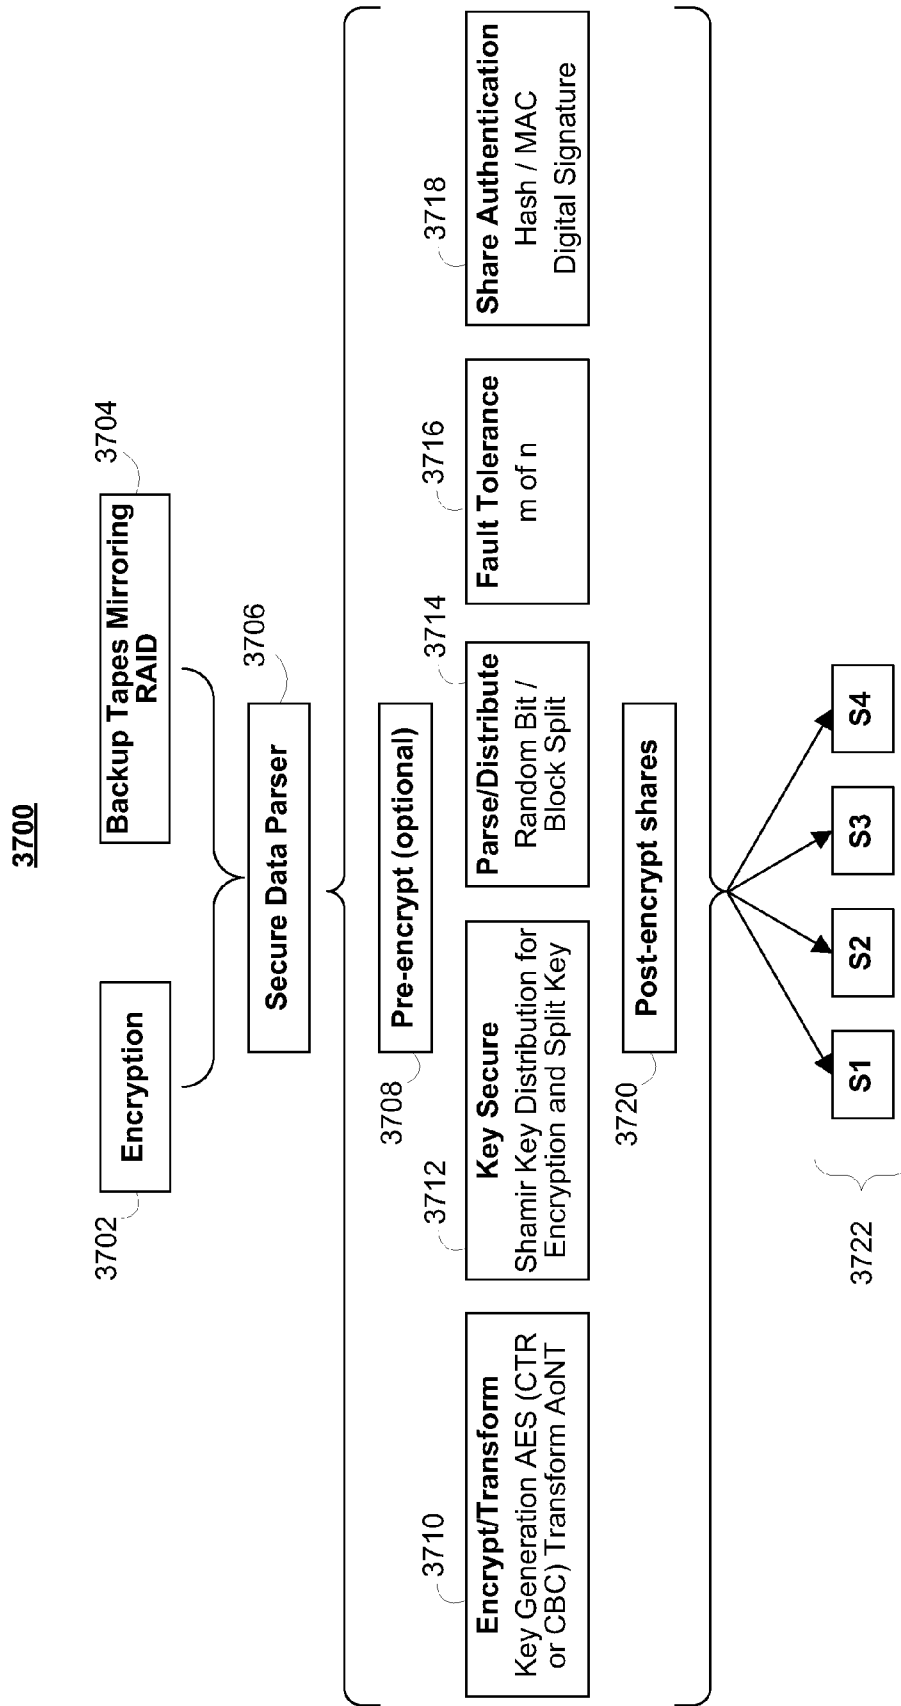


FIG. 37

3800




FIG. 38




FIG. 39




FIG. 40B




FIG. 40A

4100

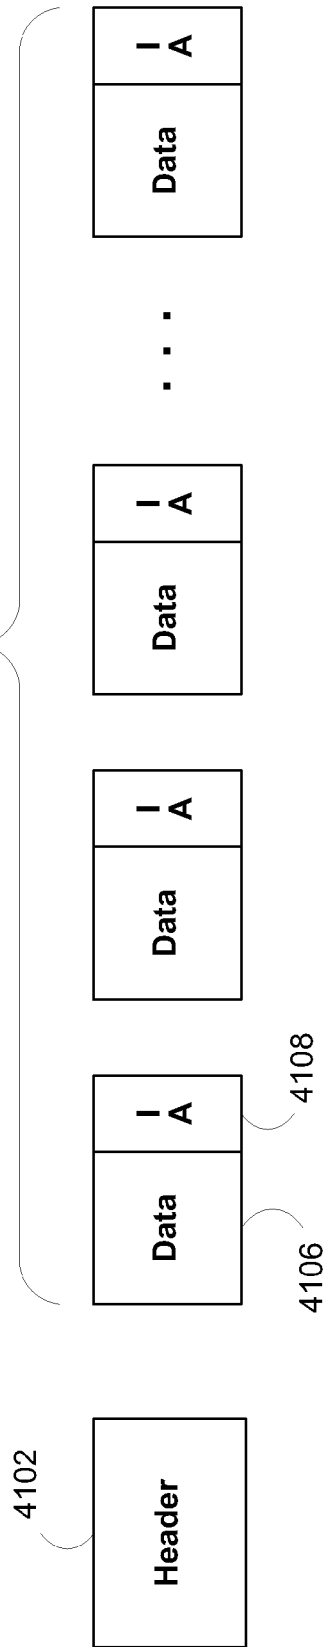


FIG. 41




FIG. 42




FIG. 43

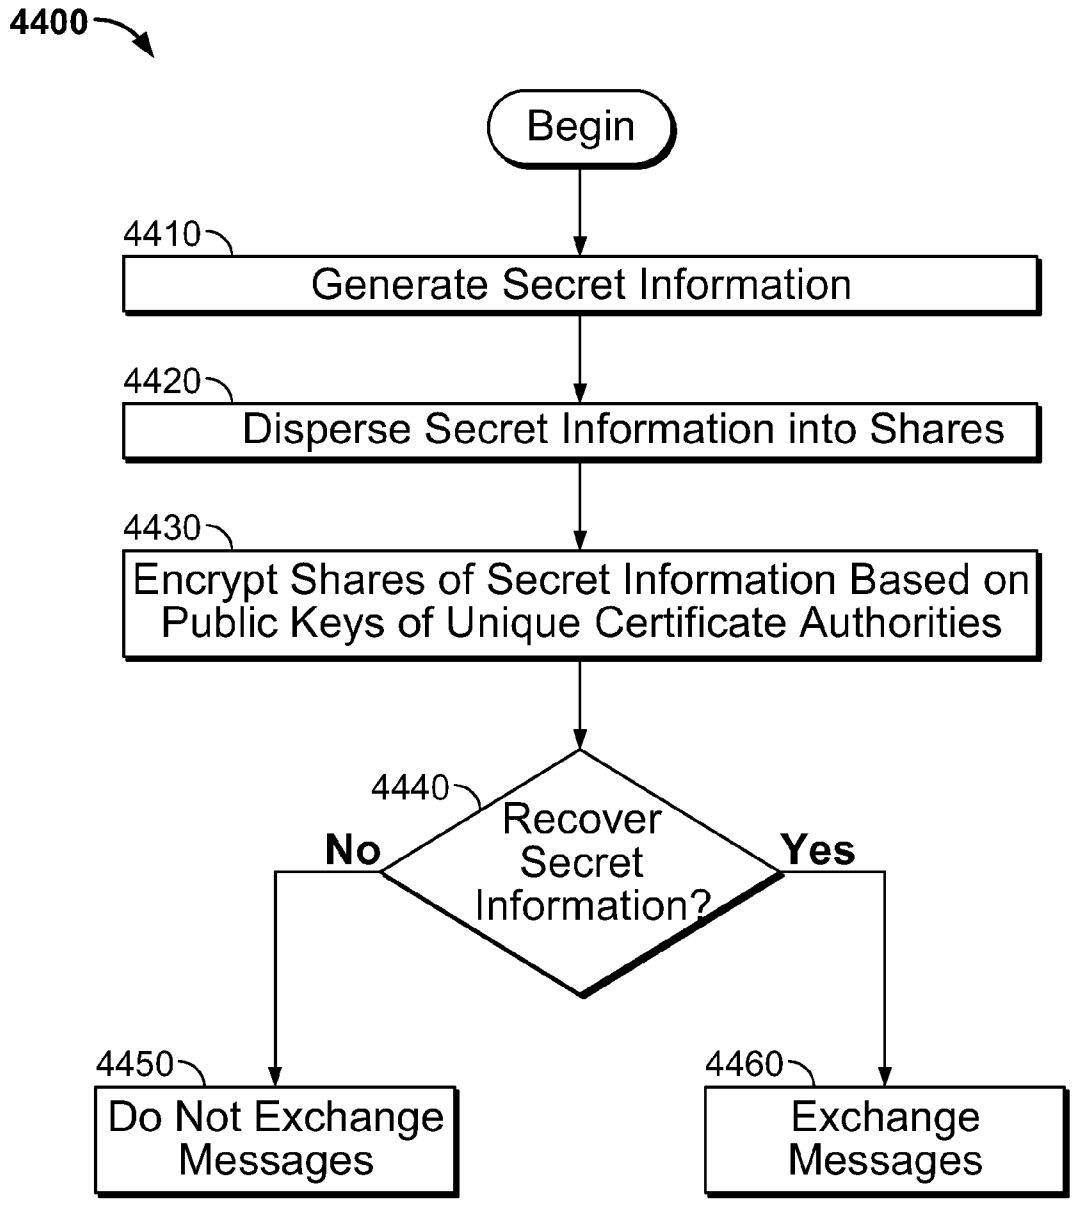


FIG. 44

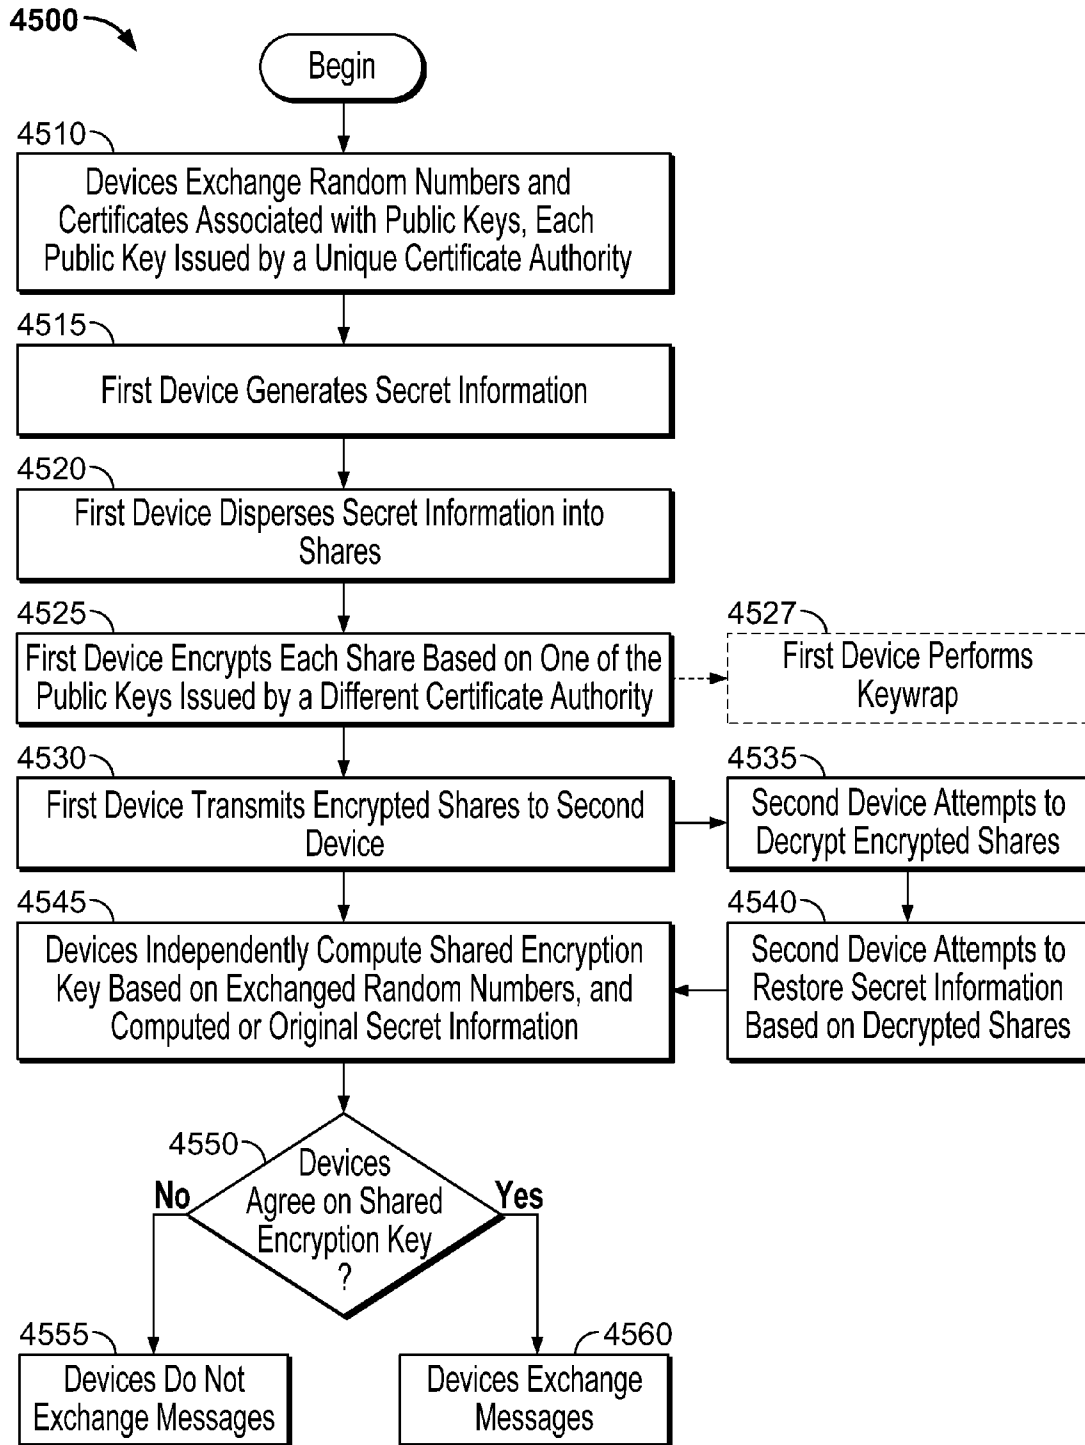



FIG. 45



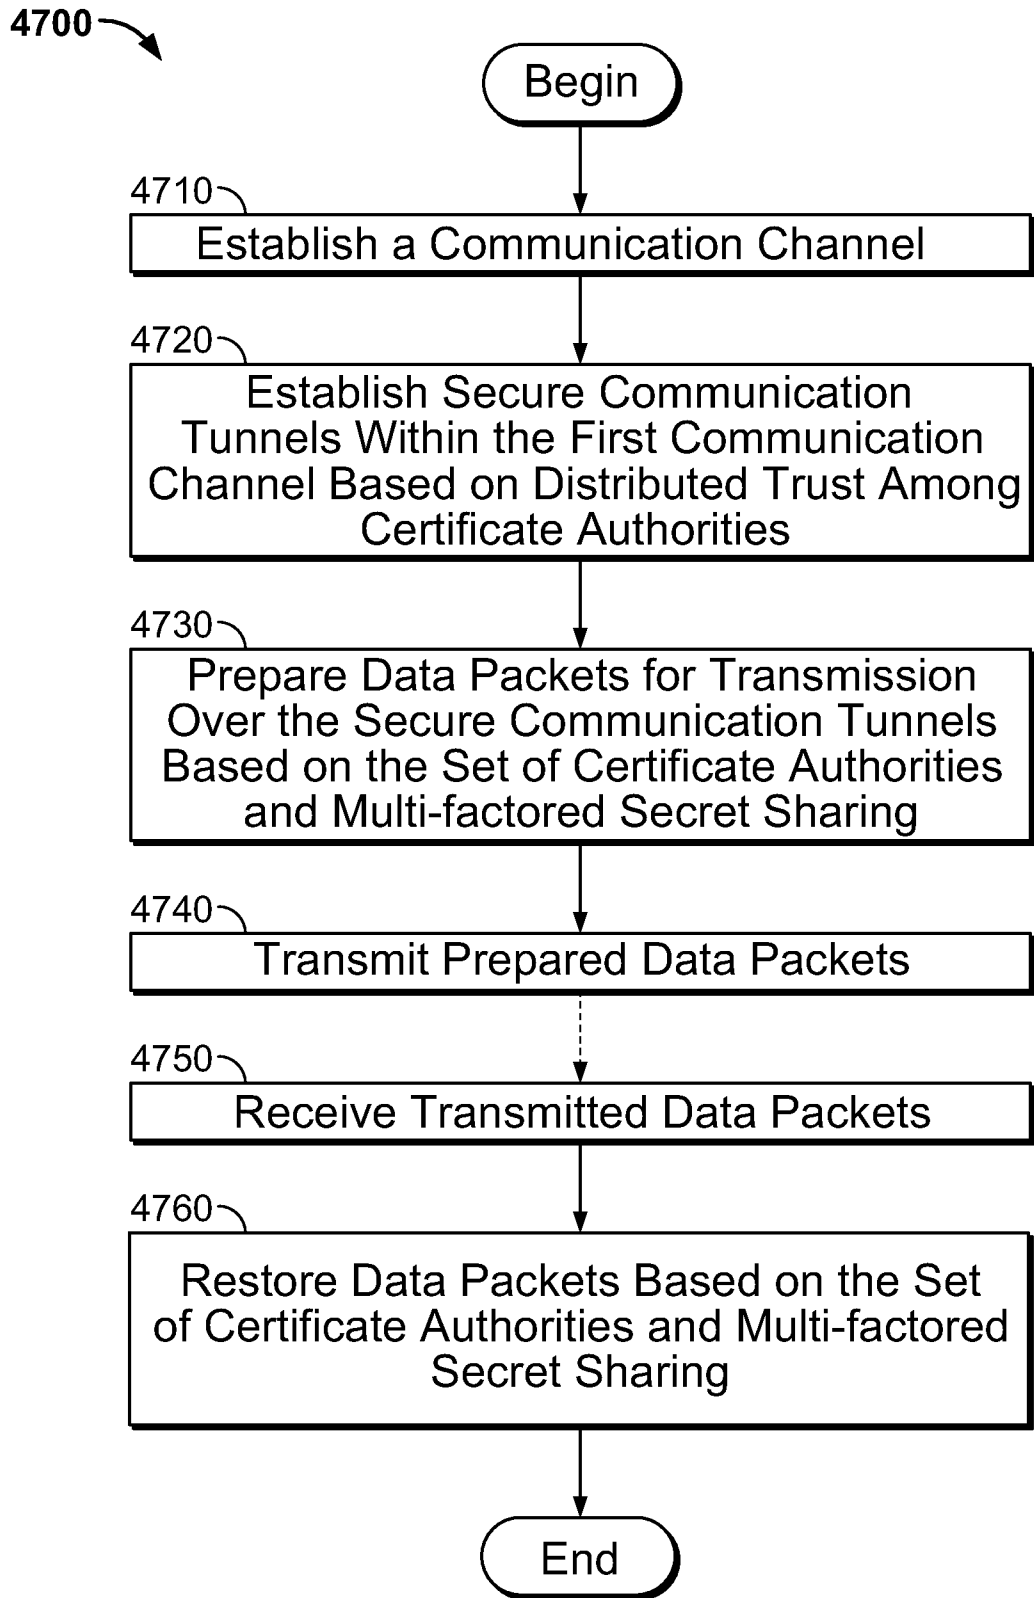


FIG. 47

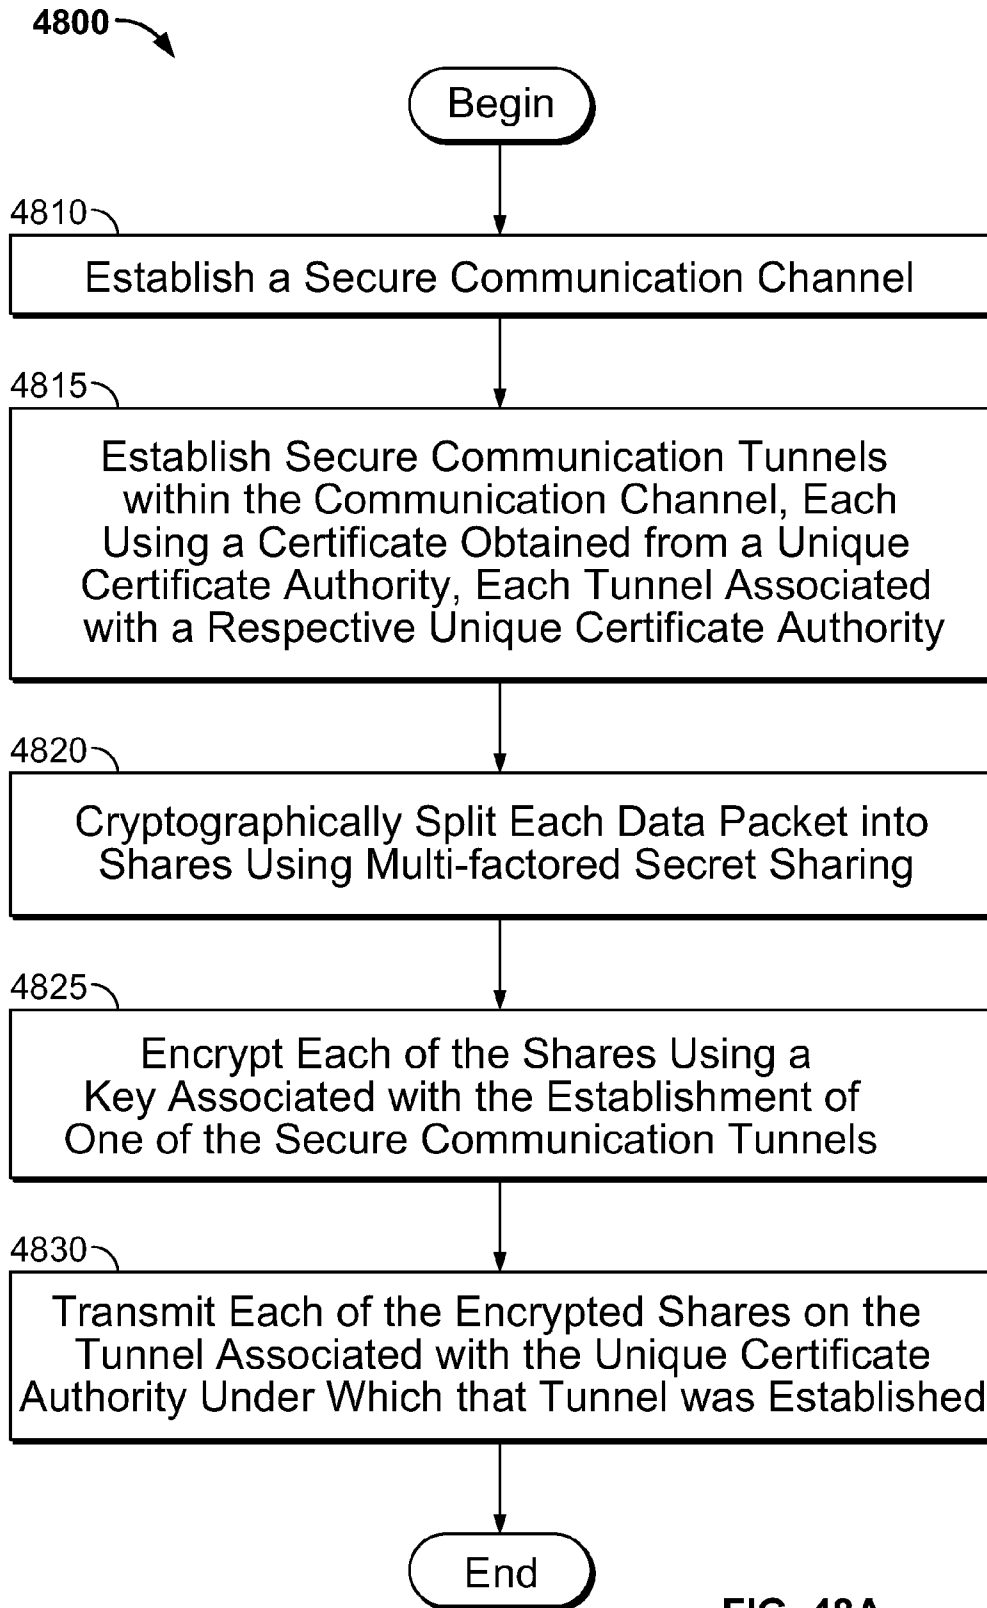


FIG. 48A

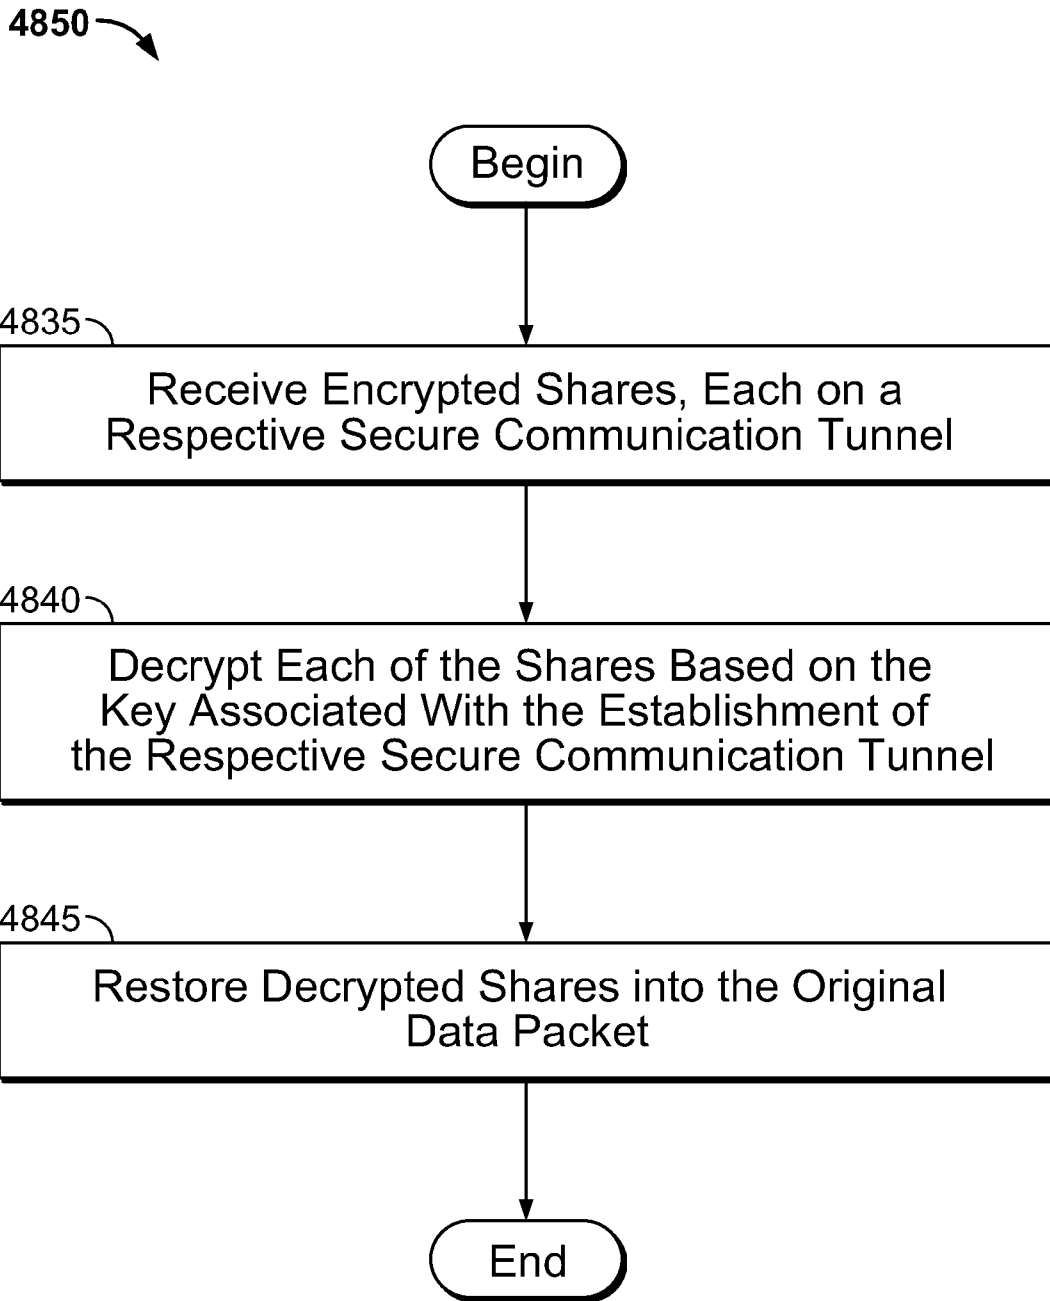


FIG. 48B




FIG. 48C




FIG. 49




FIG. 50

SYSTEMS AND METHODS FOR SECURING DATA IN MOTION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. provisional application No. 61/264,464, filed on Nov. 25, 2009, which is hereby incorporated by reference herein in its entirety.

FIELD OF THE INVENTION

The present invention relates in general to systems and methods for securing communications by distributing trust among certificate authorities. The systems and methods described herein may be used in conjunction with other systems and methods described in commonly-owned U.S. Pat. No. 7,391,865 and commonly-owned U.S. patent application Ser. No. 11/258,839, filed Oct. 25, 2005, Ser. No. 11/602,667, filed Nov. 20, 2006, Ser. No. 11/983,355, filed Nov. 7, 2007, Ser. No. 11/999,575, filed Dec. 5, 2007, Ser. No. 12/148,365, filed Apr. 18, 2008, Ser. No. 12/209,703, filed Sep. 12, 2008, Ser. No. 12/349,897, filed Jan. 7, 2009, Ser. No. 12/391,025, filed Feb. 23, 2009, all of which are hereby incorporated by reference herein in their entireties.

BACKGROUND OF THE INVENTION

In today's society, individuals and businesses conduct an ever-increasing amount of activities on and over computer systems. These computer systems, including proprietary and non-proprietary computer networks, are often storing, archiving, and transmitting all types of sensitive information. Thus, an ever-increasing need exists for ensuring data stored and transmitted over these systems cannot be read or otherwise compromised.

One solution is to secure the data using keys of a certificate authority. Certificate authorities may be run by trusted third-party organizations or companies that issue digital certificates, such as, for example, VeriSign, Baltimore, Entrust, or the like. The digital certificate certifies the ownership of a public key by the named subject of the certificate. This allows others to rely upon signatures or assertions made by the private key that corresponds to the public key that is certified. Requests for a digital certificate may be made through digital certificate protocols, such as, for example, PKCS10. In response to a request, the certificate authority will issue a certificate in a number of differing protocols, such as, for example, PKCS7. Messages may be exchanged between devices based on the issued certificates.

If the certificate authority is compromised, then the security of the system may be lost for each user for whom the certificate authority is certifying a link between a public key and an identity. For example, an attacker may compromise a certificate authority by inducing that certificate authority to issue a certificate that falsely claims to represent an entity. The attacker would have the private key associated with the certificate authority's certificate. The attacker could then use this certificate to send digitally signed messages to a user, and trick that user into believing that the message was from the trusted entity. The user may respond to the digitally signed messages, which the attacker may decrypt using the private key. Accordingly, the trust that the user placed in the certificate authority may be compromised.

SUMMARY OF THE INVENTION

Based on the foregoing, a need exists to provide a secure proxy service that includes a system that secures communications by distributing trust among a set of certificate authorities.

Accordingly, the present invention provides two approaches to distributing trust among a set of certificate authorities. Both approaches are equally secure. In each approach, a secure data parser may be integrated with any suitable encryption technology. It will be understood that in some embodiments the secure proxy service may be implemented by integrating a secure data parser with the full Transport Layer Security ("TLS") protocol, with the Secure Sockets Layer (SSL) protocol, with SSL and full TLS, or implementing the secure data parser without the use of SSL and/or full TLS. In addition, it will be understood that in some embodiments the secure proxy service may be implemented in conjunction with any suitable protocol that makes use of certificate authorities to ensure the confidentiality, integrity, and authenticity of exchanged messages.

Accordingly, one approach of the present invention provides methods and systems in which the secure data parser is used to distribute trust in a set of certificate authorities during initial negotiation (e.g., the key establishment phase) of a connection between devices. The certificate authorities may be unique in that the certificates issued by each have different public and private key pairs. This offers the assurance that if some (but fewer than a quorum) of the certificate authorities have been compromised, the connection can still be established, and messages may be exchanged without disrupting the confidentiality or integrity of the communication.

One aspect of this approach is to provide methods and systems for computing shared encryption keys. This computation of shared encryption keys may be part of a key establishment phase of secure communications between devices. Secret information may be generated, and public keys may be obtained from unique certificate authorities. The secret information may be dispersed into any number of shares of secret information. Each share of secret information may be encrypted based on a public key of a certificate associated with a different one of the unique certificate authorities. Optionally, each of the shares of secret information may be encrypted based on a keywrap. The keywrap may be based on a workgroup key. In some embodiments, the shares may be recombined, and data may be transmitted based on the recombined shares.

In some embodiments, a set of random numbers may be generated. A first shared encryption key may be computed based on the set of random numbers and the original secret information. A second shared encryption key may be computed based on the set of random numbers and the recombined shares. Data may be transmitted based on the first and second shared encryption keys. In some embodiments, the first and second shared encryption keys may be compared. A determination may be made whether to transmit data based on this comparison, and data may be transmitted based on this determination.

Another approach of the present invention is to provide methods and systems in which the secure data parser is used to pre-process packets of data. The pre-processed data may then be dispersed into shares. A set of tunnels may be established within a communication channel using certificates issued by unique certificate authorities, keys developed during this establishment may be used to encrypt data for each of the tunnels, and the individual shares of data may be transmitted on each of the tunnels. Thus, in the second approach

trust may be distributed among a set of certificate authorities in the structure of the communication channel itself.

One aspect of this approach is to provide methods and systems for securing data in motion. The data in motion may include original data packets. A secure communication channel may be established. Any number of secure communication tunnels may be established within the secure communication channel based on a distributed trust among unique certificate authorities. In some embodiments, each of the secure communication tunnels may be established using a certificate issued by a different one of unique certificate authorities. Each of the original data packets may be prepared for transmission over the secure communication tunnels based on the distributed trust among the set of certificate authorities and multi-factored secret sharing. In some embodiments, each of the original data packets may be dispersed into a plurality of shares based on multi-factored secret sharing. Optionally, this dispersing may be based on an M of N cryptosplit. The shares may be encrypted based on a key associated with the establishment of a different one of the secure communication tunnels. In some embodiments, each of the encrypted shares may be transmitted over the secure communication tunnel used to encrypt that share.

In some embodiments, each secure communication tunnels may be established based on the certificate associated with a different one of the unique certificate authorities. In some embodiments, each secure communication tunnel may be associated with a certificate issued by one of the unique certificate authorities under which the secure communication tunnels were established. For example, there may be a one-to-one correspondence between a certificate of one of the unique certificate authorities and the secure communication tunnels. In some embodiments, these associations may be dynamic. In some embodiments, the shares may be restored by recombining at least a quorum of the shares.

In some embodiments, each of the encrypted shares may be received on a respective one of the secure communication tunnels. Each of the shares may be decrypted based on the keys associated with the establishment of the secure communication tunnels. The original data packets may be restored based on, for example the multi-factored secret sharing.

In some embodiments, a certificate authority hierarchy may be generated. The certificate authority hierarchy may include a set of root certificate authorities, a set of minor certificate authorities, or both. The set of certificate authorities may include the set of root certificate authorities, the set of minor certificate authorities, or both from the generated certificate authority hierarchy.

In some embodiments, each of the N secure communication tunnels may be established over different physical transport mediums. In some embodiments, at least one of the physical transport mediums may fail, but the original data packets may be restored without a loss of data integrity. In some embodiments, a portion of the shares is designated for transmission over at least one of the failed physical transport mediums, but some of the physical transport mediums may be operational. In such embodiments, additional secure communication tunnels may be established within at least one of the operational physical transport mediums. The portion of the shares designated for transmission over the at least one failed physical transport mediums may be transmitted over the additional secure communication tunnels.

In some embodiments, the two approaches may be combined in any suitable way. For example, any number of the secure communication tunnels in the second approach may be established using the key establishment techniques of the first approach.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described in more detail below in connection with the attached drawings, which are meant to illustrate and not to limit the invention, and in which:

FIG. 1 illustrates a block diagram of a cryptographic system, according to aspects of an embodiment of the invention;

FIG. 2 illustrates a block diagram of the trust engine of FIG. 1, according to aspects of an embodiment of the invention;

FIG. 3 illustrates a block diagram of the transaction engine of FIG. 2, according to aspects of an embodiment of the invention;

FIG. 4 illustrates a block diagram of the depository of FIG. 2, according to aspects of an embodiment of the invention;

FIG. 5 illustrates a block diagram of the authentication engine of FIG. 2, according to aspects of an embodiment of the invention;

FIG. 6 illustrates a block diagram of the cryptographic engine of FIG. 2, according to aspects of an embodiment of the invention;

FIG. 7 illustrates a block diagram of a depository system, according to aspects of another embodiment of the invention;

FIG. 8 illustrates a flow chart of a data splitting process according to aspects of an embodiment of the invention;

FIG. 9, Panel A illustrates a data flow of an enrollment process according to aspects of an embodiment of the invention;

FIG. 9, Panel B illustrates a flow chart of an interoperability process according to aspects of an embodiment of the invention;

FIG. 10 illustrates a data flow of an authentication process according to aspects of an embodiment of the invention;

FIG. 11 illustrates a data flow of a signing process according to aspects of an embodiment of the invention;

FIG. 12 illustrates a data flow and an encryption/decryption process according to aspects and yet another embodiment of the invention;

FIG. 13 illustrates a simplified block diagram of a trust engine system according to aspects of another embodiment of the invention;

FIG. 14 illustrates a simplified block diagram of a trust engine system according to aspects of another embodiment of the invention;

FIG. 15 illustrates a block diagram of the redundancy module of FIG. 14, according to aspects of an embodiment of the invention;

FIG. 16 illustrates a process for evaluating authentications according to one aspect of the invention;

FIG. 17 illustrates a process for assigning a value to an authentication according to one aspect as shown in FIG. 16 of the invention;

FIG. 18 illustrates a process for performing trust arbitrage in an aspect of the invention as shown in FIG. 17; and

FIG. 19 illustrates a sample transaction between a user and a vendor according to aspects of an embodiment of the invention where an initial web based contact leads to a sales contract signed by both parties.

FIG. 20 illustrates a sample user system with a cryptographic service provider module which provides security functions to a user system.

FIG. 21 illustrates a process for parsing, splitting and/or separating data with encryption and storage of the encryption master key with the data.

FIG. 22 illustrates a process for parsing, splitting and/or separating data with encryption and storing the encryption master key separately from the data.

FIG. 23 illustrates the intermediary key process for parsing, splitting and/or separating data with encryption and storage of the encryption master key with the data.

FIG. 24 illustrates the intermediary key process for parsing, splitting and/or separating data with encryption and storing the encryption master key separately from the data.

FIG. 25 illustrates utilization of the cryptographic methods and systems of the present invention with a small working group.

FIG. 26 is a block diagram of an illustrative physical token security system employing the secure data parser in accordance with one embodiment of the present invention.

FIG. 27 is a block diagram of an illustrative arrangement in which the secure data parser is integrated into a system in accordance with one embodiment of the present invention.

FIG. 28 is a block diagram of an illustrative data in motion system in accordance with one embodiment of the present invention.

FIG. 29 is a block diagram of another illustrative data in motion system in accordance with one embodiment of the present invention.

FIG. 30-32 are block diagrams of an illustrative system having the secure data parser integrated in accordance with one embodiment of the present invention.

FIG. 33 is a process flow diagram of an illustrative process for parsing and splitting data in accordance with one embodiment of the present invention.

FIG. 34 is a process flow diagram of an illustrative process for restoring portions of data into original data in accordance with one embodiment of the present invention.

FIG. 35 is a process flow diagram of an illustrative process for splitting data at the bit level in accordance with one embodiment of the present invention.

FIG. 36 is a process flow diagram of illustrative steps and features in accordance with one embodiment of the present invention.

FIG. 37 is a process flow diagram of illustrative steps and features in accordance with one embodiment of the present invention.

FIG. 38 is a simplified block diagram of the storage of key and data components within shares in accordance with one embodiment of the present invention.

FIG. 39 is a simplified block diagram of the storage of key and data components within shares using a workgroup key in accordance with one embodiment of the present invention.

FIGS. 40A and 40B are simplified and illustrative process flow diagrams for header generation and data splitting for data in motion in accordance with one embodiment of the present invention.

FIG. 41 is a simplified block diagram of an illustrative share format in accordance with one embodiment of the present invention.

FIG. 42 is a simplified and an illustrative hierarchy of certificate authorities in accordance with one embodiment of the present invention.

FIGS. 43-47, 48A, and 48B are process flow diagrams of illustrative steps and features for a secure proxy service in accordance with one embodiment of the present invention.

FIG. 48C is a simplified block diagram of a secure proxy service that distributes trust among a set of certificate authorities in the structure of communication channels in accordance with one embodiment of the present invention.

FIGS. 49 and 50 are process flow diagrams of illustrated steps and features for a secure proxy service that distributes trust among a set of certificate authorities in the structure of communication channels in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

One aspect of the present invention is to provide a cryptographic system where one or more secure servers, or a trust engine, stores cryptographic keys and user authentication data. Users access the functionality of conventional cryptographic systems through network access to the trust engine, however, the trust engine does not release actual keys and other authentication data and therefore, the keys and data remain secure. This server-centric storage of keys and authentication data provides for user-independent security, portability, availability, and straightforwardness.

Because users can be confident in, or trust, the cryptographic system to perform user and document authentication and other cryptographic functions, a wide variety of functionality may be incorporated into the system. For example, the trust engine provider can ensure against agreement repudiation by, for example, authenticating the agreement participants, digitally signing the agreement on behalf of or for the participants, and storing a record of the agreement digitally signed by each participant. In addition, the cryptographic system may monitor agreements and determine to apply varying degrees of authentication, based on, for example, price, user, vendor, geographic location, place of use, or the like.

To facilitate a complete understanding of the invention, the remainder of the detailed description describes the invention with reference to the figures, wherein like elements are referenced with like numerals throughout.

FIG. 1 illustrates a block diagram of a cryptographic system 100, according to aspects of an embodiment of the invention. As shown in FIG. 1, the cryptographic system 100 includes a user system 105, a trust engine 110, a certificate authority 115, and a vendor system 120, communicating through a communication link 125.

According to one embodiment of the invention, the user system 105 comprises a conventional general-purpose computer having one or more microprocessors, such as, for example, an Intel-based processor. Moreover, the user system 105 includes an appropriate operating system, such as, for example, an operating system capable of including graphics or windows, such as Windows, Unix, Linux, or the like. As shown in FIG. 1, the user system 105 may include a biometric device 107. The biometric device 107 may advantageously capture a user's biometric and transfer the captured biometric to the trust engine 110. According to one embodiment of the invention, the biometric device may advantageously comprise a device having attributes and features similar to those disclosed in U.S. patent application Ser. No. 08/926,277, filed on Sep. 5, 1997, entitled "RELIEF OBJECT IMAGE GENERATOR," U.S. patent application Ser. No. 09/558,634, filed on Apr. 26, 2000, entitled "IMAGING DEVICE FOR A RELIEF OBJECT AND SYSTEM AND METHOD OF USING THE IMAGE DEVICE," U.S. patent application Ser. No. 09/435,011, filed on Nov. 5, 1999, entitled "RELIEF OBJECT SENSOR ADAPTOR," and U.S. patent application Ser. No. 09/477,943, filed on Jan. 5, 2000, entitled "PLANAR OPTICAL IMAGE SENSOR AND SYSTEM FOR GENERATING AN ELECTRONIC IMAGE OF A RELIEF OBJECT FOR FINGERPRINT READING," all of which are owned by the instant assignee, and all of which are hereby incorporated by reference herein.

In addition, the user system 105 may connect to the communication link 125 through a conventional service provider, such as, for example, a dial up, digital subscriber line (DSL), cable modem, fiber connection, or the like. According to another embodiment, the user system 105 connects the communication link 125 through network connectivity such as,

for example, a local or wide area network. According to one embodiment, the operating system includes a TCP/IP stack that handles all incoming and outgoing message traffic passed over the communication link 125.

Although the user system 105 is disclosed with reference to the foregoing embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives 5 embodiments of the user system 105, including almost any computing device capable of sending or receiving information from another computer system. For example, the user system 105 may include, but is not limited to, a computer workstation, an interactive television, an interactive kiosk, a personal mobile computing device, such as a digital assistant, mobile phone, laptop, or the like, personal networking equip- 10 ment, such as a home router, a network storage device ("NAS"), personal hotspot, or the like, or a wireless communications device, a smartcard, an embedded computing device, or the like, which can interact with the communication link 125. In such alternative systems, the operating systems will likely differ and be adapted for the particular device. However, according to one embodiment, the operating systems advantageously continue to provide the appropriate 15 communications protocols needed to establish communication with the communication link 125.

FIG. 1 illustrates the trust engine 110. According to one embodiment, the trust engine 110 comprises one or more secure servers for accessing and storing sensitive information, which may be any type or form of data, such as, but not limited to text, audio, video, user authentication data and public and private cryptographic keys. According to one embodiment, the authentication data includes data designed to uniquely identify a user of the cryptographic system 100. For example, the authentication data may include a user identification number, one or more biometrics, and a series of 35 questions and answers generated by the trust engine 110 or the user, but answered initially by the user at enrollment. The foregoing questions may include demographic data, such as place of birth, address, anniversary, or the like, personal data, such as mother's maiden name, favorite ice cream, or the like, or other data designed to uniquely identify the user. The trust engine 110 compares a user's authentication data associated with a current transaction, to the authentication data provided at an earlier time, such as, for example, during enrollment. The trust engine 110 may advantageously require the user to produce the authentication data at the time of each transaction, or, the trust engine 110 may advantageously allow the user to periodically produce authentication data, such as at the beginning of a string of transactions or the logging onto a particular vendor website.

According to the embodiment where the user produces biometric data, the user provides a physical characteristic, such as, but not limited to, facial scan, hand scan, ear scan, iris scan, retinal scan, vascular pattern, DNA, a fingerprint, writing or speech, to the biometric device 107. The biometric device advantageously produces an electronic pattern, or biometric, of the physical characteristic. The electronic pattern is transferred through the user system 105 to the trust engine 110 for either enrollment or authentication purposes.

Once the user produces the appropriate authentication data and the trust engine 110 determines a positive match between that authentication data (current authentication data) and the authentication data provided at the time of enrollment (enrollment authentication data), the trust engine 110 provides the user with complete cryptographic functionality. For example, the properly authenticated user may advantageously employ the trust engine 110 to perform hashing, digitally signing,

encrypting and decrypting (often together referred to only as encrypting), creating or distributing digital certificates, and the like. However, the private cryptographic keys used in the cryptographic functions will not be available outside the trust engine 110, thereby ensuring the integrity of the cryptographic keys.

According to one embodiment, the trust engine 110 generates and stores cryptographic keys. According to another embodiment, at least one cryptographic key is associated with each user. Moreover, when the cryptographic keys include public-key technology, each private key associated with a user is generated within, and not released from, the trust engine 110. Thus, so long as the user has access to the trust engine 110, the user may perform cryptographic functions using his or her private or public key. Such remote access advantageously allows users to remain completely mobile and access cryptographic functionality through practically any Internet connection, such as cellular and satellite phones, kiosks, laptops, hotel rooms and the like.

According to another embodiment, the trust engine 110 performs the cryptographic functionality using a key pair generated for the trust engine 110. According to this embodiment, the trust engine 110 first authenticates the user, and after the user has properly produced authentication data matching the enrollment authentication data, the trust engine 110 uses its own cryptographic key pair to perform cryptographic functions on behalf of the authenticated user.

A skilled artisan will recognize from the disclosure herein that the cryptographic keys may advantageously include some or all of symmetric keys, public keys, and private keys. In addition, a skilled artisan will recognize from the disclosure herein that the foregoing keys may be implemented with a wide number of algorithms available from commercial technologies, such as, for example, RSA, ELGAMAL, or the like.

FIG. 1 also illustrates the certificate authority 115. According to one embodiment, the certificate authority 115 may advantageously comprise a trusted third-party organization or company that issues digital certificates, such as, for example, VeriSign, Baltimore, Entrust, or the like. The trust engine 110 may advantageously transmit requests for digital certificates, through one or more conventional digital certificate protocols, such as, for example, PKCS10, to the certificate authority 115. In response, the certificate authority 115 will issue a digital certificate in one or more of a number of differing protocols, such as, for example, PKCS7. According to one embodiment of the invention, the trust engine 110 requests digital certificates from several or all of the prominent certificate authorities 115 such that the trust engine 110 has access to a digital certificate corresponding to the certificate standard of any requesting party.

According to another embodiment, the trust engine 110 internally performs certificate issuances. In this embodiment, the trust engine 110 may access a certificate system for generating certificates and/or may internally generate certificates when they are requested, such as, for example, at the time of key generation or in the certificate standard requested at the time of the request. The trust engine 110 will be disclosed in greater detail below.

FIG. 1 also illustrates the vendor system 120. According to one embodiment, the vendor system 120 advantageously comprises a Web server. Typical Web servers generally serve content over the Internet using one of several internet markup languages or document format standards, such as the Hypertext Markup Language (HTML) or the Extensible Markup Language (XML). The Web server accepts requests from browsers like Netscape and Internet Explorer and then returns the appropriate electronic documents. A number of server or

client-side technologies can be used to increase the power of the Web server beyond its ability to deliver standard electronic documents. For example, these technologies include Common Gateway Interface (CGI) scripts, SSL security, and Active Server Pages (ASPs). The vendor system **120** may advantageously provide electronic content relating to commercial, personal, educational, or other transactions.

Although the vendor system **120** is disclosed with reference to the foregoing embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein that the vendor system **120** may advantageously comprise any of the devices described with reference to the user system **105** or combination thereof.

FIG. **1** also illustrates the communication link **125** connecting the user system **105**, the trust engine **110**, the certificate authority **115**, and the vendor system **120**. According to one embodiment, the communication link **125** preferably comprises the Internet. The Internet, as used throughout this disclosure is a global network of computers. The structure of the Internet, which is well known to those of ordinary skill in the art, includes a network backbone with networks branching from the backbone. These branches, in turn, have networks branching from them, and so on. Routers move information packets between network levels, and then from network to network, until the packet reaches the neighborhood of its destination. From the destination, the destination network's host directs the information packet to the appropriate terminal, or node. In one advantageous embodiment, the Internet routing hubs comprise domain name system (DNS) servers using Transmission Control Protocol/Internet Protocol (TCP/IP) as is well known in the art. The routing hubs connect to one or more other routing hubs via high-speed communication links.

One popular part of the Internet is the World Wide Web. The World Wide Web contains different computers, which store documents capable of displaying graphical and textual information. The computers that provide information on the World Wide Web are typically called "websites." A website is defined by an Internet address that has an associated electronic page. The electronic page can be identified by a Uniform Resource Locator (URL). Generally, an electronic page is a document that organizes the presentation of text, graphical images, audio, video, and so forth.

Although the communication link **125** is disclosed in terms of its preferred embodiment, one of ordinary skill in the art will recognize from the disclosure herein that the communication link **125** may include a wide range of interactive communications links. For example, the communication link **125** may include interactive television networks, telephone networks, wireless data transmission systems, two-way cable systems, customized private or public computer networks, interactive kiosk networks, automatic teller machine networks, direct links, satellite or cellular networks, and the like.

FIG. **2** illustrates a block diagram of the trust engine **110** according to aspects of an embodiment of the invention. As shown in FIG. **2**, the trust engine **110** includes a transaction engine **205**, a depository **210**, an authentication engine **215**, and a cryptographic engine **220**. According to one embodiment of the invention, the trust engine **110** also includes mass storage **225**. As further shown in FIG. **2**, the transaction engine **205** communicates with the depository **210**, the authentication engine **215**, and the cryptographic engine **220**, along with the mass storage **225**. In addition, the depository **210** communicates with the authentication engine **215**, the cryptographic engine **220**, and the mass storage **225**. Moreover, the authentication engine **215** communicates with

the cryptographic engine **220**. According to one embodiment of the invention, some or all of the foregoing communications may advantageously comprise the transmission of XML documents to IP addresses that correspond to the receiving device. As mentioned in the foregoing, XML documents advantageously allow designers to create their own customized document tags, enabling the definition, transmission, validation, and interpretation of data between applications and between organizations. Moreover, some or all of the foregoing communications may include conventional SSL technologies.

According to one embodiment, the transaction engine **205** comprises a data routing device, such as a conventional Web server available from Netscape, Microsoft, Apache, or the like. For example, the Web server may advantageously receive incoming data from the communication link **125**. According to one embodiment of the invention, the incoming data is addressed to a front-end security system for the trust engine **110**. For example, the front-end security system may advantageously include a firewall, an intrusion detection system searching for known attack profiles, and/or a virus scanner. After clearing the front-end security system, the data is received by the transaction engine **205** and routed to one of the depository **210**, the authentication engine **215**, the cryptographic engine **220**, and the mass storage **225**. In addition, the transaction engine **205** monitors incoming data from the authentication engine **215** and cryptographic engine **220**, and routes the data to particular systems through the communication link **125**. For example, the transaction engine **205** may advantageously route data to the user system **105**, the certificate authority **115**, or the vendor system **120**.

According to one embodiment, the data is routed using conventional HTTP routing techniques, such as, for example, employing URLs or Uniform Resource Indicators (URIs). URIs are similar to URLs, however, URIs typically indicate the source of files or actions, such as, for example, executables, scripts, and the like. Therefore, according to the one embodiment, the user system **105**, the certificate authority **115**, the vendor system **120**, and the components of the trust engine **110**, advantageously include sufficient data within communication URLs or URIs for the transaction engine **205** to properly route data throughout the cryptographic system.

Although the data routing is disclosed with reference to its preferred embodiment, a skilled artisan will recognize a wide number of possible data routing solutions or strategies. For example, XML or other data packets may advantageously be unpacked and recognized by their format, content, or the like, such that the transaction engine **205** may properly route data throughout the trust engine **110**. Moreover, a skilled artisan will recognize that the data routing may advantageously be adapted to the data transfer protocols conforming to particular network systems, such as, for example, when the communication link **125** comprises a local network.

According to yet another embodiment of the invention, the transaction engine **205** includes conventional SSL encryption technologies, such that the foregoing systems may authenticate themselves, and vice-versa, with transaction engine **205**, during particular communications. As will be used throughout this disclosure, the term "1/2 SSL" refers to communications where a server but not necessarily the client, is SSL authenticated, and the term "FULL SSL" refers to communications where the client and the server are SSL authenticated. When the instant disclosure uses the term "SSL", the communication may comprise 1/2 or FULL SSL.

As the transaction engine **205** routes data to the various components of the cryptographic system **100**, the transaction

11

engine 205 may advantageously create an audit trail. According to one embodiment, the audit trail includes a record of at least the type and format of data routed by the transaction engine 205 throughout the cryptographic system 100. Such audit data may advantageously be stored in the mass storage 225.

FIG. 2 also illustrates the depository 210. According to one embodiment, the depository 210 comprises one or more data storage facilities, such as, for example, a directory server, a database server, or the like. As shown in FIG. 2, the depository 210 stores cryptographic keys and enrollment authentication data. The cryptographic keys may advantageously correspond to the trust engine 110 or to users of the cryptographic system 100, such as the user or vendor. The enrollment authentication data may advantageously include data designed to uniquely identify a user, such as, user ID, passwords, answers to questions, biometric data, or the like. This enrollment authentication data may advantageously be acquired at enrollment of a user or another alternative later time. For example, the trust engine 110 may include periodic or other renewal or reissue of enrollment authentication data.

According to one embodiment, the communication from the transaction engine 205 to and from the authentication engine 215 and the cryptographic engine 220 comprises secure communication, such as, for example conventional SSL technology. In addition, as mentioned in the foregoing, the data of the communications to and from the depository 210 may be transferred using URLs, URIs, HTTP or XML documents, with any of the foregoing advantageously having data requests and formats embedded therein.

As mentioned above, the depository 210 may advantageously comprise a plurality of secure data storage facilities. In such an embodiment, the secure data storage facilities may be configured such that a compromise of the security in one individual data storage facility will not compromise the cryptographic keys or the authentication data stored therein. For example, according to this embodiment, the cryptographic keys and the authentication data are mathematically operated on so as to statistically and substantially randomize the data stored in each data storage facility. According to one embodiment, the randomization of the data of an individual data storage facility renders that data undecipherable. Thus, compromise of an individual data storage facility produces only a randomized undecipherable number and does not compromise the security of any cryptographic keys or the authentication data as a whole.

FIG. 2 also illustrates the trust engine 110 including the authentication engine 215. According to one embodiment, the authentication engine 215 comprises a data comparator configured to compare data from the transaction engine 205 with data from the depository 210. For example, during authentication, a user supplies current authentication data to the trust engine 110 such that the transaction engine 205 receives the current authentication data. As mentioned in the foregoing, the transaction engine 205 recognizes the data requests, preferably in the URL or URI, and routes the authentication data to the authentication engine 215. Moreover, upon request, the depository 210 forwards enrollment authentication data corresponding to the user to the authentication engine 215. Thus, the authentication engine 215 has both the current authentication data and the enrollment authentication data for comparison.

According to one embodiment, the communications to the authentication engine comprise secure communications, such as, for example, SSL technology. Additionally, security can be provided within the trust engine 110 components, such as, for example, super-encryption using public key technologies.

12

For example, according to one embodiment, the user encrypts the current authentication data with the public key of the authentication engine 215. In addition, the depository 210 also encrypts the enrollment authentication data with the public key of the authentication engine 215. In this way, only the authentication engine's private key can be used to decrypt the transmissions.

As shown in FIG. 2, the trust engine 110 also includes the cryptographic engine 220. According to one embodiment, the cryptographic engine comprises a cryptographic handling module, configured to advantageously provide conventional cryptographic functions, such as, for example, public-key infrastructure (PKI) functionality. For example, the cryptographic engine 220 may advantageously issue public and private keys for users of the cryptographic system 100. In this manner, the cryptographic keys are generated at the cryptographic engine 220 and forwarded to the depository 210 such that at least the private cryptographic keys are not available outside of the trust engine 110. According to another embodiment, the cryptographic engine 220 randomizes and splits at least the private cryptographic key data, thereby storing only the randomized split data. Similar to the splitting of the enrollment authentication data, the splitting process ensures the stored keys are not available outside the cryptographic engine 220. According to another embodiment, the functions of the cryptographic engine can be combined with and performed by the authentication engine 215.

According to one embodiment, communications to and from the cryptographic engine include secure communications, such as SSL technology. In addition, XML documents may advantageously be employed to transfer data and/or make cryptographic function requests.

FIG. 2 also illustrates the trust engine 110 having the mass storage 225. As mentioned in the foregoing, the transaction engine 205 keeps data corresponding to an audit trail and stores such data in the mass storage 225. Similarly, according to one embodiment of the invention, the depository 210 keeps data corresponding to an audit trail and stores such data in the mass storage device 225. The depository audit trail data is similar to that of the transaction engine 205 in that the audit trail data comprises a record of the requests received by the depository 210 and the response thereof. In addition, the mass storage 225 may be used to store digital certificates having the public key of a user contained therein.

Although the trust engine 110 is disclosed with reference to its preferred and alternative embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize in the disclosure herein, a wide number of alternatives for the trust engine 110. For example, the trust engine 110, may advantageously perform only authentication, or alternatively, only some or all of the cryptographic functions, such as data encryption and decryption. According to such embodiments, one of the authentication engine 215 and the cryptographic engine 220 may advantageously be removed, thereby creating a more straightforward design for the trust engine 110. In addition, the cryptographic engine 220 may also communicate with a certificate authority such that the certificate authority is embodied within the trust engine 110. According to yet another embodiment, the trust engine 110 may advantageously perform authentication and one or more cryptographic functions, such as, for example, digital signing.

FIG. 3 illustrates a block diagram of the transaction engine 205 of FIG. 2, according to aspects of an embodiment of the invention. According to this embodiment, the transaction engine 205 comprises an operating system 305 having a handling thread and a listening thread. The operating system 305

may advantageously be similar to those found in conventional high volume servers, such as, for example, Web servers available from Apache. The listening thread monitors the incoming communication from one of the communication link **125**, the authentication engine **215**, and the cryptographic engine **220** for incoming data flow. The handling thread recognizes particular data structures of the incoming data flow, such as, for example, the foregoing data structures, thereby routing the incoming data to one of the communication link **125**, the depository **210**, the authentication engine **215**, the cryptographic engine **220**, or the mass storage **225**. As shown in FIG. 3, the incoming and outgoing data may advantageously be secured through, for example, SSL technology.

FIG. 4 illustrates a block diagram of the depository **210** of FIG. 2 according to aspects of an embodiment of the invention. According to this embodiment, the depository **210** comprises one or more lightweight directory access protocol (LDAP) servers. LDAP directory servers are available from a wide variety of manufacturers such as Netscape, ISO, and others. FIG. 4 also shows that the directory server preferably stores data **405** corresponding to the cryptographic keys and data **410** corresponding to the enrollment authentication data. According to one embodiment, the depository **210** comprises a single logical memory structure indexing authentication data and cryptographic key data to a unique user ID. The single logical memory structure preferably includes mechanisms to ensure a high degree of trust, or security, in the data stored therein. For example, the physical location of the depository **210** may advantageously include a wide number of conventional security measures, such as limited employee access, modern surveillance systems, and the like. In addition to, or in lieu of, the physical securities, the computer system or server may advantageously include software solutions to protect the stored data. For example, the depository **210** may advantageously create and store data **415** corresponding to an audit trail of actions taken. In addition, the incoming and outgoing communications may advantageously be encrypted with public key encryption coupled with conventional SSL technologies.

According to another embodiment, the depository **210** may comprise distinct and physically separated data storage facilities, as disclosed further with reference to FIG. 7.

FIG. 5 illustrates a block diagram of the authentication engine **215** of FIG. 2 according to aspects of an embodiment of the invention. Similar to the transaction engine **205** of FIG. 3, the authentication engine **215** comprises an operating system **505** having at least a listening and a handling thread of a modified version of a conventional Web server, such as, for example, Web servers available from Apache. As shown in FIG. 5, the authentication engine **215** includes access to at least one private key **510**. The private key **510** may advantageously be used for example, to decrypt data from the transaction engine **205** or the depository **210**, which was encrypted with a corresponding public key of the authentication engine **215**.

FIG. 5 also illustrates the authentication engine **215** comprising a comparator **515**, a data splitting module **520**, and a data assembling module **525**. According to the preferred embodiment of the invention, the comparator **515** includes technology capable of comparing potentially complex patterns related to the foregoing biometric authentication data. The technology may include hardware, software, or combined solutions for pattern comparisons, such as, for example, those representing finger print patterns or voice patterns. In addition, according to one embodiment, the comparator **515** of the authentication engine **215** may advantageously compare conventional hashes of documents in order to render a

comparison result. According to one embodiment of the invention, the comparator **515** includes the application of heuristics **530** to the comparison. The heuristics **530** may advantageously address circumstances surrounding an authentication attempt, such as, for example, the time of day, IP address or subnet mask, purchasing profile, email address, processor serial number or ID, or the like.

Moreover, the nature of biometric data comparisons may result in varying degrees of confidence being produced from the matching of current biometric authentication data to enrollment data. For example, unlike a traditional password which may only return a positive or negative match, a fingerprint may be determined to be a partial match, e.g. a 90% match, a 75% match, or a 10% match, rather than simply being correct or incorrect. Other biometric identifiers such as voice print analysis or face recognition may share this property of probabilistic authentication, rather than absolute authentication.

When working with such probabilistic authentication or in other cases where an authentication is considered less than absolutely reliable, it is desirable to apply the heuristics **530** to determine whether the level of confidence in the authentication provided is sufficiently high to authenticate the transaction which is being made.

It will sometimes be the case that the transaction at issue is a relatively low value transaction where it is acceptable to be authenticated to a lower level of confidence. This could include a transaction which has a low dollar value associated with it (e.g., a \$10 purchase) or a transaction with low risk (e.g., admission to a members-only web site).

Conversely, for authenticating other transactions, it may be desirable to require a high degree of confidence in the authentication before allowing the transaction to proceed. Such transactions may include transactions of large dollar value (e.g., signing a multi-million dollar supply contract) or transaction with a high risk if an improper authentication occurs (e.g., remotely logging onto a government computer).

The use of the heuristics **530** in combination with confidence levels and transactions values may be used as will be described below to allow the comparator to provide a dynamic context-sensitive authentication system.

According to another embodiment of the invention, the comparator **515** may advantageously track authentication attempts for a particular transaction. For example, when a transaction fails, the trust engine **110** may request the user to re-enter his or her current authentication data. The comparator **515** of the authentication engine **215** may advantageously employ an attempt limiter **535** to limit the number of authentication attempts, thereby prohibiting brute-force attempts to impersonate a user's authentication data. According to one embodiment, the attempt limiter **535** comprises a software module monitoring transactions for repeating authentication attempts and, for example, limiting the authentication attempts for a given transaction to three. Thus, the attempt limiter **535** will limit an automated attempt to impersonate an individual's authentication data to, for example, simply three "guesses." Upon three failures, the attempt limiter **535** may advantageously deny additional authentication attempts. Such denial may advantageously be implemented through, for example, the comparator **515** returning a negative result regardless of the current authentication data being transmitted. On the other hand, the transaction engine **205** may advantageously block any additional authentication attempts pertaining to a transaction in which three attempts have previously failed.

The authentication engine **215** also includes the data splitting module **520** and the data assembling module **525**. The

15

data splitting module 520 advantageously comprises a software, hardware, or combination module having the ability to mathematically operate on various data so as to substantially randomize and split the data into portions. According to one embodiment, original data is not recreatable from an individual portion. The data assembling module 525 advantageously comprises a software, hardware, or combination module configured to mathematically operate on the foregoing substantially randomized portions, such that the combination thereof provides the original deciphered data. According to one embodiment, the authentication engine 215 employs the data splitting module 520 to randomize and split enrollment authentication data into portions, and employs the data assembling module 525 to reassemble the portions into usable enrollment authentication data.

FIG. 6 illustrates a block diagram of the cryptographic engine 220 of the trust engine 200 of FIG. 2 according to aspects of one embodiment of the invention. Similar to the transaction engine 205 of FIG. 3, the cryptographic engine 220 comprises an operating system 605 having at least a listening and a handling thread of a modified version of a conventional Web server, such as, for example, Web servers available from Apache. As shown in FIG. 6, the cryptographic engine 220 comprises a data splitting module 610 and a data assembling module 620 that function similar to those of FIG. 5. However, according to one embodiment, the data splitting module 610 and the data assembling module 620 process cryptographic key data, as opposed to the foregoing enrollment authentication data. Although, a skilled artisan will recognize from the disclosure herein that the data splitting module 910 and the data splitting module 620 may be combined with those of the authentication engine 215.

The cryptographic engine 220 also comprises a cryptographic handling module 625 configured to perform one, some or all of a wide number of cryptographic functions. According to one embodiment, the cryptographic handling module 625 may comprise software modules or programs, hardware, or both. According to another embodiment, the cryptographic handling module 625 may perform data comparisons, data parsing, data splitting, data separating, data hashing, data encryption or decryption, digital signature verification or creation, digital certificate generation, storage, or requests, cryptographic key generation, or the like. Moreover, a skilled artisan will recognize from the disclosure herein that the cryptographic handling module 825 may advantageously comprise a public-key infrastructure, such as Pretty Good Privacy (PGP), an RSA-based public-key system, or a wide number of alternative key management systems. In addition, the cryptographic handling module 625 may perform public-key encryption, symmetric-key encryption, or both. In addition to the foregoing, the cryptographic handling module 625 may include one or more computer programs or modules, hardware, or both, for implementing seamless, transparent, interoperability functions.

A skilled artisan will also recognize from the disclosure herein that the cryptographic functionality may include a wide number or variety of functions generally relating to cryptographic key management systems.

FIG. 7 illustrates a simplified block diagram of a depository system 700 according to aspects of an embodiment of the invention. As shown in FIG. 7, the depository system 700 advantageously comprises multiple data storage facilities, for example, data storage facilities D1, D2, D3, and D4. However, it is readily understood by those of ordinary skill in the art that the depository system may have only one data storage facility. According to one embodiment of the invention, each of the data storage facilities D1 through D4 may advantageously

16

comprise some or all of the elements disclosed with reference to the depository 210 of FIG. 4. Similar to the depository 210, the data storage facilities D1 through D4 communicate with the transaction engine 205, the authentication engine 215, and the cryptographic engine 220, preferably through conventional SSL. Communication links transferring, for example, XML documents. Communications from the transaction engine 205 may advantageously include requests for data, wherein the request is advantageously broadcast to the IP address of each data storage facility D1 through D4. On the other hand, the transaction engine 205 may broadcast requests to particular data storage facilities based on a wide number of criteria, such as, for example, response time, server loads, maintenance schedules, or the like.

In response to requests for data from the transaction engine 205, the depository system 700 advantageously forwards stored data to the authentication engine 215 and the cryptographic engine 220. The respective data assembling modules receive the forwarded data and assemble the data into useable formats. On the other hand, communications from the authentication engine 215 and the cryptographic engine 220 to the data storage facilities D1 through D4 may include the transmission of sensitive data to be stored. For example, according to one embodiment, the authentication engine 215 and the cryptographic engine 220 may advantageously employ their respective data splitting modules to divide sensitive data into undecipherable portions, and then transmit one or more undecipherable portions of the sensitive data to a particular data storage facility.

According to one embodiment, each data storage facility, D1 through D4, comprises a separate and independent storage system, such as, for example, a directory server. According to another embodiment of the invention, the depository system 700 comprises multiple geographically separated independent data storage systems. By distributing the sensitive data into distinct and independent storage facilities D1 through D4, some or all of which may be advantageously geographically separated, the depository system 700 provides redundancy along with additional security measures. For example, according to one embodiment, only data from two of the multiple data storage facilities, D1 through D4, are needed to decipher and reassemble the sensitive data. Thus, as many as two of the four data storage facilities D1 through D4 may be inoperative due to maintenance, system failure, power failure, or the like, without affecting the functionality of the trust engine 110. In addition, because, according to one embodiment, the data stored in each data storage facility is randomized and undecipherable, compromise of any individual data storage facility does not necessarily compromise the sensitive data. Moreover, in the embodiment having geographical separation of the data storage facilities, a compromise of multiple geographically remote facilities becomes increasingly difficult. In fact, even a rogue employee will be greatly challenged to subvert the needed multiple independent geographically remote data storage facilities.

Although the depository system 700 is disclosed with reference to its preferred and alternative embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for the depository system 700. For example, the depository system 700 may comprise one, two or more data storage facilities. In addition, sensitive data may be mathematically operated such that portions from two or more data storage facilities are needed to reassemble and decipher the sensitive data.

As mentioned in the foregoing, the authentication engine 215 and the cryptographic engine 220 each include a data splitting module 520 and 610, respectively, for splitting any type or form of sensitive data, such as, for example, text, audio, video, the authentication data and the cryptographic key data. FIG. 8 illustrates a flowchart of a data splitting process 800 performed by the data splitting module according to aspects of an embodiment of the invention. As shown in FIG. 8, the data splitting process 800 begins at step 805 when sensitive data "S" is received by the data splitting module of the authentication engine 215 or the cryptographic engine 220. Preferably, in step 810, the data splitting module then generates a substantially random number, value, or string or set of bits, "A." For example, the random number A may be generated in a wide number of varying conventional techniques available to one of ordinary skill in the art, for producing high quality random numbers suitable for use in cryptographic applications. In addition, according to one embodiment, the random number A comprises a bit length which may be any suitable length, such as shorter, longer or equal to the bit length of the sensitive data, S.

In addition, in step 820 the data splitting process 800 generates another statistically random number "C." According to the preferred embodiment, the generation of the statistically random numbers A and C may advantageously be done in parallel. The data splitting module then combines the numbers A and C with the sensitive data S such that new numbers "B" and "D" are generated. For example, number B may comprise the binary combination of A XOR S and number D may comprise the binary combination of C XOR S. The XOR function, or the "exclusive-or" function, is well known to those of ordinary skill in the art. The foregoing combinations preferably occur in steps 825 and 830, respectively, and, according to one embodiment, the foregoing combinations also occur in parallel. The data splitting process 800 then proceeds to step 835 where the random numbers A and C and the numbers B and D are paired such that none of the pairings contain sufficient data, by themselves, to reorganize and decipher the original sensitive data S. For example, the numbers may be paired as follows: AC, AD, BC, and BD. According to one embodiment, each of the foregoing pairings is distributed to one of the depositories D1 through D4 of FIG. 7. According to another embodiment, each of the foregoing pairings is randomly distributed to one of the depositories D1 through D4. For example, during a first data splitting process 800, the pairing AC may be sent to depository D2, through, for example, a random selection of D2's IP address. Then, during a second data splitting process 800, the pairing AC may be sent to depository D4, through, for example, a random selection of D4's IP address. In addition, the pairings may all be stored on one depository, and may be stored in separate locations on said depository.

Based on the foregoing, the data splitting process 800 advantageously places portions of the sensitive data in each of the four data storage facilities D1 through D4, such that no single data storage facility D1 through D4 includes sufficient encrypted data to recreate the original sensitive data S. As mentioned in the foregoing, such randomization of the data into individually unusable encrypted portions increases security and provides for maintained trust in the data even if one of the data storage facilities, D1 through D4, is compromised.

Although the data splitting process 800 is disclosed with reference to its preferred embodiment, the invention is not intended to be limited thereby. Rather a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for the data splitting process 800. For example, the data splitting process may advantageously split the data into

two numbers, for example, random number A and number B and, randomly distribute A and B through two data storage facilities. Moreover, the data splitting process 800 may advantageously split the data among a wide number of data storage facilities through generation of additional random numbers. The data may be split into any desired, selected, predetermined, or randomly assigned size unit, including but not limited to, a bit, bits, bytes, kilobytes, megabytes or larger, or any combination or sequence of sizes. In addition, varying the sizes of the data units resulting from the splitting process may render the data more difficult to restore to a useable form, thereby increasing security of sensitive data. It is readily apparent to those of ordinary skill in the art that the split data unit sizes may be a wide variety of data unit sizes or patterns of sizes or combinations of sizes. For example, the data unit sizes may be selected or predetermined to be all of the same size, a fixed set of different sizes, a combination of sizes, or randomly generates sizes. Similarly, the data units may be distributed into one or more shares according to a fixed or predetermined data unit size, a pattern or combination of data unit sizes, or a randomly generated data unit size or sizes per share.

As mentioned in the foregoing, in order to recreate the sensitive data S, the data portions need to be derandomized and reorganized. This process may advantageously occur in the data assembling modules, 525 and 620, of the authentication engine 215 and the cryptographic engine 220, respectively. The data assembling module, for example, data assembly module 525, receives data portions from the data storage facilities D1 through D4, and reassembles the data into useable form. For example, according to one embodiment where the data splitting module 520 employed the data splitting process 800 of FIG. 8, the data assembling module 525 uses data portions from at least two of the data storage facilities D1 through D4 to recreate the sensitive data S. For example, the pairings of AC, AD, BC, and BD, were distributed such that any two provide one of A and B, or, C and D. Noting that $S=A \text{ XOR } B$ or $S=C \text{ XOR } D$ indicates that when the data assembling module receives one of A and B, or, C and D, the data assembling module 525 can advantageously reassemble the sensitive data S. Thus, the data assembling module 525 may assemble the sensitive data S, when, for example, it receives data portions from at least the first two of the data storage facilities D1 through D4 to respond to an assemble request by the trust engine 110.

Based on the above data splitting and assembling processes, the sensitive data S exists in usable format only in a limited area of the trust engine 110. For example, when the sensitive data S includes enrollment authentication data, usable, nonrandomized enrollment authentication data is available only in the authentication engine 215. Likewise, when the sensitive data S includes private cryptographic key data, usable, nonrandomized private cryptographic key data is available only in the cryptographic engine 220.

Although the data splitting and assembling processes are disclosed with reference to their preferred embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for splitting and reassembling the sensitive data S. For example, public-key encryption may be used to further secure the data at the data storage facilities D1 through D4. In addition, it is readily apparent to those of ordinary skill in the art that the data splitting module described herein is also a separate and distinct embodiment of the present invention that may be incorporated into, combined with or otherwise made part of any pre-existing computer systems, software suites, database, or combinations thereof,

or other embodiments of the present invention, such as the trust engine, authentication engine, and transaction engine disclosed and described herein.

FIG. 9A illustrates a data flow of an enrollment process 900 according to aspects of an embodiment of the invention. As shown in FIG. 9A, the enrollment process 900 begins at step 905 when a user desires to enroll with the trust engine 110 of the cryptographic system 100. According to this embodiment, the user system 105 advantageously includes a client-side applet, such as a Java-based, that queries the user to enter enrollment data, such as demographic data and enrollment authentication data. According to one embodiment, the enrollment authentication data includes user ID, password(s), biometric(s), or the like. According to one embodiment, during the querying process, the client-side applet preferably communicates with the trust engine 110 to ensure that a chosen user ID is unique. When the user ID is nonunique, the trust engine 110 may advantageously suggest a unique user ID. The client-side applet gathers the enrollment data and transmits the enrollment data, for example, through an XML document, to the trust engine 110, and in particular, to the transaction engine 205. According to one embodiment, the transmission is encoded with the public key of the authentication engine 215.

According to one embodiment, the user performs a single enrollment during step 905 of the enrollment process 900. For example, the user enrolls himself or herself as a particular person, such as Joe User. When Joe User desires to enroll as Joe User, CEO of Mega Corp., then according to this embodiment, Joe User enrolls a second time, receives a second unique user ID and the trust engine 110 does not associate the two identities. According to another embodiment of the invention, the enrollment process 900 provides for multiple user identities for a single user ID. Thus, in the above example, the trust engine 110 will advantageously associate the two identities of Joe User. As will be understood by a skilled artisan from the disclosure herein, a user may have many identities, for example, Joe User the head of household, Joe User the member of the Charitable Foundations, and the like. Even though the user may have multiple identities, according to this embodiment, the trust engine 110 preferably stores only one set of enrollment data. Moreover, users may advantageously add, edit/update, or delete identities as they are needed.

Although the enrollment process 900 is disclosed with reference to its preferred embodiment, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for gathering of enrollment data, and in particular, enrollment authentication data. For example, the applet may be common object model (COM) based applet or the like.

On the other hand, the enrollment process may include graded enrollment. For example, at a lowest level of enrollment, the user may enroll over the communication link 125 without producing documentation as to his or her identity. According to an increased level of enrollment, the user enrolls using a trusted third party, such as a digital notary. For example, and the user may appear in person to the trusted third party, produce credentials such as a birth certificate, driver's license, military ID, or the like, and the trusted third party may advantageously include, for example, their digital signature in enrollment submission. The trusted third party may include an actual notary, a government agency, such as the Post Office or Department of Motor Vehicles, a human resources person in a large company enrolling an employee, or the like. A skilled artisan will understand from the disclo-

sure herein that a wide number of varying levels of enrollment may occur during the enrollment process 900.

After receiving the enrollment authentication data, at step 915, the transaction engine 205, using conventional FULL SSL technology forwards the enrollment authentication data to the authentication engine 215. In step 920, the authentication engine 215 decrypts the enrollment authentication data using the private key of the authentication engine 215. In addition, the authentication engine 215 employs the data splitting module to mathematically operate on the enrollment authentication data so as to split the data into at least two independently undecipherable, randomized, numbers. As mentioned in the foregoing, at least two numbers may comprise a statistically random number and a binary XORed number. In step 925, the authentication engine 215 forwards each portion of the randomized numbers to one of the data storage facilities D1 through D4. As mentioned in the foregoing, the authentication engine 215 may also advantageously randomize which portions are transferred to which depositories.

Often during the enrollment process 900, the user will also desire to have a digital certificate issued such that he or she may receive encrypted documents from others outside the cryptographic system 100. As mentioned in the foregoing, the certificate authority 115 generally issues digital certificates according to one or more of several conventional standards. Generally, the digital certificate includes a public key of the user or system, which is known to everyone.

Whether the user requests a digital certificate at enrollment, or at another time, the request is transferred through the trust engine 110 to the authentication engine 215. According to one embodiment, the request includes an XML document having, for example, the proper name of the user. According to step 935, the authentication engine 215 transfers the request to the cryptographic engine 220 instructing the cryptographic engine 220 to generate a cryptographic key or key pair.

Upon request, at step 935, the cryptographic engine 220 generates at least one cryptographic key. According to one embodiment, the cryptographic handling module 625 generates a key pair, where one key is used as a private key, and one is used as a public key. The cryptographic engine 220 stores the private key and, according to one embodiment, a copy of the public key. In step 945, the cryptographic engine 220 transmits a request for a digital certificate to the transaction engine 205. According to one embodiment, the request advantageously includes a standardized request, such as PKCS10, embedded in, for example, an XML document. The request for a digital certificate may advantageously correspond to one or more certificate authorities and the one or more standard formats the certificate authorities require.

In step 950 the transaction engine 205 forwards this request to the certificate authority 115, who, in step 955, returns a digital certificate. The return digital certificate may advantageously be in a standardized format, such as PKCS7, or in a proprietary format of one or more of the certificate authorities 115. In step 960, the digital certificate is received by the transaction engine 205, and a copy is forwarded to the user and a copy is stored with the trust engine 110. The trust engine 110 stores a copy of the certificate such that the trust engine 110 will not need to rely on the availability of the certificate authority 115. For example, when the user desires to send a digital certificate, or a third party requests the user's digital certificate, the request for the digital certificate is typically sent to the certificate authority 115. However, if the certificate

authority 115 is conducting maintenance or has been victim of a failure or security compromise, the digital certificate may not be available.

At any time after issuing the cryptographic keys, the cryptographic engine 220 may advantageously employ the data splitting process 800 described above such that the cryptographic keys are split into independently undecipherable randomized numbers. Similar to the authentication data, at step 965 the cryptographic engine 220 transfers the randomized numbers to the data storage facilities D1 through D4.

A skilled artisan will recognize from the disclosure herein that the user may request a digital certificate anytime after enrollment. Moreover, the communications between systems may advantageously include FULL SSL or public-key encryption technologies. Moreover, the enrollment process may issue multiple digital certificates from multiple certificate authorities, including one or more proprietary certificate authorities internal or external to the trust engine 110.

As disclosed in steps 935 through 960, one embodiment of the invention includes the request for a certificate that is eventually stored on the trust engine 110. Because, according to one embodiment, the cryptographic handling module 625 issues the keys used by the trust engine 110, each certificate corresponds to a private key. Therefore, the trust engine 110 may advantageously provide for interoperability through monitoring the certificates owned by, or associated with, a user. For example, when the cryptographic engine 220 receives a request for a cryptographic function, the cryptographic handling module 625 may investigate the certificates owned by the requesting user to determine whether the user owns a private key matching the attributes of the request. When such a certificate exists, the cryptographic handling module 625 may use the certificate or the public or private keys associated therewith, to perform the requested function. When such a certificate does not exist, the cryptographic handling module 625 may advantageously and transparently perform a number of actions to attempt to remedy the lack of an appropriate key. For example, FIG. 9B illustrates a flowchart of an interoperability process 970, which according to aspects of an embodiment of the invention, discloses the foregoing steps to ensure the cryptographic handling module 625 performs cryptographic functions using appropriate keys.

As shown in FIG. 9B, the interoperability process 970 begins with step 972 where the cryptographic handling module 925 determines the type of certificate desired. According to one embodiment of the invention, the type of certificate may advantageously be specified in the request for cryptographic functions, or other data provided by the requestor. According to another embodiment, the certificate type may be ascertained by the data format of the request. For example, the cryptographic handling module 925 may advantageously recognize the request corresponds to a particular type.

According to one embodiment, the certificate type may include one or more algorithm standards, for example, RSA, ELGAMAL, or the like. In addition, the certificate type may include one or more key types, such as symmetric keys, public keys, strong encryption keys such as 256 bit keys, less secure keys, or the like. Moreover, the certificate type may include upgrades or replacements of one or more of the foregoing algorithm standards or keys, one or more message or data formats, one or more data encapsulation or encoding schemes, such as Base 32 or Base 64. The certificate type may also include compatibility with one or more third-party cryptographic applications or interfaces, one or more communication protocols, or one or more certificate standards or protocols. A skilled artisan will recognize from the disclosure

herein that other differences may exist in certificate types, and translations to and from those differences may be implemented as disclosed herein.

Once the cryptographic handling module 625 determines the certificate type, the interoperability process 970 proceeds to step 974, and determines whether the user owns a certificate matching the type determined in step 974. When the user owns a matching certificate, for example, the trust engine 110 has access to the matching certificate through, for example, prior storage thereof, the cryptographic handling module 825 knows that a matching private key is also stored within the trust engine 110. For example, the matching private key may be stored within the depository 210 or depository system 700. The cryptographic handling module 625 may advantageously request the matching private key be assembled from, for example, the depository 210, and then in step 976, use the matching private key to perform cryptographic actions or functions. For example, as mentioned in the foregoing, the cryptographic handling module 625 may advantageously perform hashing, hash comparisons, data encryption or decryption, digital signature verification or creation, or the like.

When the user does not own a matching certificate, the interoperability process 970 proceeds to step 978 where the cryptographic handling module 625 determines whether the user owns a cross-certified certificate. According to one embodiment, cross-certification between certificate authorities occurs when a first certificate authority determines to trust certificates from a second certificate authority. In other words, the first certificate authority determines that certificates from the second certificate authority meets certain quality standards, and therefore, may be "certified" as equivalent to the first certificate authority's own certificates. Cross-certification becomes more complex when the certificate authorities issue, for example, certificates having levels of trust. For example, the first certificate authority may provide three levels of trust for a particular certificate, usually based on the degree of reliability in the enrollment process, while the second certificate authority may provide seven levels of trust. Cross-certification may advantageously track which levels and which certificates from the second certificate authority may be substituted for which levels and which certificates from the first. When the foregoing cross-certification is done officially and publicly between two certification authorities, the mapping of certificates and levels to one another is often called "chaining."

According to another embodiment of the invention, the cryptographic handling module 625 may advantageously develop cross-certifications outside those agreed upon by the certificate authorities. For example, the cryptographic handling module 625 may access a first certificate authority's certificate practice statement (CPS), or other published policy statement, and using, for example, the authentication tokens required by particular trust levels, match the first certificate authority's certificates to those of another certificate authority.

When, in step 978, the cryptographic handling module 625 determines that the user owns a cross-certified certificate, the interoperability process 970 proceeds to step 976, and performs the cryptographic action or function using the cross-certified public key, private key, or both. Alternatively, when the cryptographic handling module 625 determines that the user does not own a cross-certified certificate, the interoperability process 970 proceeds to step 980, where the cryptographic handling module 625 selects a certificate authority that issues the requested certificate type, or a certificate cross-certified thereto. In step 982, the cryptographic handling module 625 determines whether the user enrollment authen-

tication data, discussed in the foregoing, meets the authentication requirements of the chosen certificate authority. For example, if the user enrolled over a network by, for example, answering demographic and other questions, the authentication data provided may establish a lower level of trust than a user providing biometric data and appearing before a third-party, such as, for example, a notary. According to one embodiment, the foregoing authentication requirements may advantageously be provided in the chosen authentication authority's CPS.

When the user has provided the trust engine 110 with enrollment authentication data meeting the requirements of chosen certificate authority, the interoperability process 970 proceeds to step 984, where the cryptographic handling module 825 acquires the certificate from the chosen certificate authority. According to one embodiment, the cryptographic handling module 625 acquires the certificate by following steps 945 through 960 of the enrollment process 900. For example, the cryptographic handling module 625 may advantageously employ one or more public keys from one or more of the key pairs already available to the cryptographic engine 220, to request the certificate from the certificate authority. According to another embodiment, the cryptographic handling module 625 may advantageously generate one or more new key pairs, and use the public keys corresponding thereto, to request the certificate from the certificate authority.

According to another embodiment, the trust engine 110 may advantageously include one or more certificate issuing modules capable of issuing one or more certificate types. According to this embodiment, the certificate issuing module may provide the foregoing certificate. When the cryptographic handling module 625 acquires the certificate, the interoperability process 970 proceeds to step 976, and performs the cryptographic action or function using the public key, private key, or both corresponding to the acquired certificate.

When the user, in step 982, has not provided the trust engine 110 with enrollment authentication data meeting the requirements of chosen certificate authority, the cryptographic handling module 625 determines, in step 986 whether there are other certificate authorities that have different authentication requirements. For example, the cryptographic handling module 625 may look for certificate authorities having lower authentication requirements, but still issue the chosen certificates, or cross-certifications thereof.

When the foregoing certificate authority having lower requirements exists, the interoperability process 970 proceeds to step 980 and chooses that certificate authority. Alternatively, when no such certificate authority exists, in step 988, the trust engine 110 may request additional authentication tokens from the user. For example, the trust engine 110 may request new enrollment authentication data comprising, for example, biometric data. Also, the trust engine 110 may request the user appear before a trusted third party and provide appropriate authenticating credentials, such as, for example, appearing before a notary with a drivers license, social security card, bank card, birth certificate, military ID, or the like. When the trust engine 110 receives updated authentication data, the interoperability process 970 proceeds to step 984 and acquires the foregoing chosen certificate.

Through the foregoing interoperability process 970, the cryptographic handling module 625 advantageously provides seamless, transparent, translations and conversions between differing cryptographic systems. A skilled artisan will recognize from the disclosure herein, a wide number of advantages and implementations of the foregoing interoperable system. For example, the foregoing step 986 of the interoperability

process 970 may advantageously include aspects of trust arbitrage, discussed in further detail below, where the certificate authority may under special circumstances accept lower levels of cross-certification. In addition, the interoperability process 970 may include ensuring interoperability between and employment of standard certificate revocations, such as employing certificate revocation lists (CRL), online certificate status protocols (OCSP), or the like.

FIG. 10 illustrates a data flow of an authentication process 1000 according to aspects of an embodiment of the invention. According to one embodiment, the authentication process 1000 includes gathering current authentication data from a user and comparing that to the enrollment authentication data of the user. For example, the authentication process 1000 begins at step 1005 where a user desires to perform a transaction with, for example, a vendor. Such transactions may include, for example, selecting a purchase option, requesting access to a restricted area or device of the vendor system 120, or the like. At step 1010, a vendor provides the user with a transaction ID and an authentication request. The transaction ID may advantageously include a 192 bit quantity having a 32 bit timestamp concatenated with a 128 bit random quantity, or a "nonce," concatenated with a 32 bit vendor specific constant. Such a transaction ID uniquely identifies the transaction such that copycat transactions can be refused by the trust engine 110.

The authentication request may advantageously include what level of authentication is needed for a particular transaction. For example, the vendor may specify a particular level of confidence that is required for the transaction at issue. If authentication cannot be made to this level of confidence, as will be discussed below, the transaction will not occur without either further authentication by the user to raise the level of confidence, or a change in the terms of the authentication between the vendor and the server. These issues are discussed more completely below.

According to one embodiment, the transaction ID and the authentication request may be advantageously generated by a vendor-side applet or other software program. In addition, the transmission of the transaction ID and authentication data may include one or more XML documents encrypted using conventional SSL technology, such as, for example, 1/2 SSL, or, in other words vendor-side authenticated SSL.

After the user system 105 receives the transaction ID and authentication request, the user system 105 gathers the current authentication data, potentially including current biometric information, from the user. The user system 105, at step 1015, encrypts at least the current authentication data "B" and the transaction ID, with the public key of the authentication engine 215, and transfers that data to the trust engine 110. The transmission preferably comprises XML documents encrypted with at least conventional 1/2 SSL technology. In step 1020, the transaction engine 205 receives the transmission, preferably recognizes the data format or request in the URL or URI, and forwards the transmission to the authentication engine 215.

During steps 1015 and 1020, the vendor system 120, at step 1025, forwards the transaction ID and the authentication request to the trust engine 110, using the preferred FULL SSL technology. This communication may also include a vendor ID, although vendor identification may also be communicated through a non-random portion of the transaction ID. At steps 1030 and 1035, the transaction engine 205 receives the communication, creates a record in the audit trail, and generates a request for the user's enrollment authentication data to be assembled from the data storage facilities D1 through D4. At step 1040, the depository system 700 transfers the portions

of the enrollment authentication data corresponding to the user to the authentication engine 215. At step 1045, the authentication engine 215 decrypts the transmission using its private key and compares the enrollment authentication data to the current authentication data provided by the user.

The comparison of step 1045 may advantageously apply heuristical context sensitive authentication, as referred to in the forgoing, and discussed in further detail below. For example, if the biometric information received does not match perfectly, a lower confidence match results. In particular embodiments, the level of confidence of the authentication is balanced against the nature of the transaction and the desires of both the user and the vendor. Again, this is discussed in greater detail below.

At step 1050, the authentication engine 215 fills in the authentication request with the result of the comparison of step 1045. According to one embodiment of the invention, the authentication request is filled with a YES/NO or TRUE/FALSE result of the authentication process 1000. In step 1055 the filled-in authentication request is returned to the vendor for the vendor to act upon, for example, allowing the user to complete the transaction that initiated the authentication request. According to one embodiment, a confirmation message is passed to the user.

Based on the foregoing, the authentication process 1000 advantageously keeps sensitive data secure and produces results configured to maintain the integrity of the sensitive data. For example, the sensitive data is assembled only inside the authentication engine 215. For example, the enrollment authentication data is undecipherable until it is assembled in the authentication engine 215 by the data assembling module, and the current authentication data is undecipherable until it is unwrapped by the conventional SSL technology and the private key of the authentication engine 215. Moreover, the authentication result transmitted to the vendor does not include the sensitive data, and the user may not even know whether he or she produced valid authentication data.

Although the authentication process 1000 is disclosed with reference to its preferred and alternative embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for the authentication process 1000. For example, the vendor may advantageously be replaced by almost any requesting application, even those residing with the user system 105. For example, a client application, such as Microsoft Word, may use an application program interface (API) or a cryptographic API (CAPI) to request authentication before unlocking a document. Alternatively, a mail server, a network, a cellular phone, a personal or mobile computing device, a workstation, or the like, may all make authentication requests that can be filled by the authentication process 1000. In fact, after providing the foregoing trusted authentication process 1000, the requesting application or device may provide access to or use of a wide number of electronic or computer devices or systems.

Moreover, the authentication process 1000 may employ a wide number of alternative procedures in the event of authentication failure. For example, authentication failure may maintain the same transaction ID and request that the user reenter his or her current authentication data. As mentioned in the foregoing, use of the same transaction ID allows the comparator of the authentication engine 215 to monitor and limit the number of authentication attempts for a particular transaction, thereby creating a more secure cryptographic system 100.

In addition, the authentication process 1000 may be advantageously be employed to develop elegant single sign-on

solutions, such as, unlocking a sensitive data vault. For example, successful or positive authentication may provide the authenticated user the ability to automatically access any number of passwords for an almost limitless number of systems and applications. For example, authentication of a user may provide the user access to password, login, financial credentials, or the like, associated with multiple online vendors, a local area network, various personal computing devices, Internet service providers, auction providers, investment brokerages, or the like. By employing a sensitive data vault, users may choose truly large and random passwords because they no longer need to remember them through association. Rather, the authentication process 1000 provides access thereto. For example, a user may choose a random alphanumeric string that is twenty plus digits in length rather than something associated with a memorable data, name, etc.

According to one embodiment, a sensitive data vault associated with a given user may advantageously be stored in the data storage facilities of the depository 210, or split and stored in the depository system 700. According to this embodiment, after positive user authentication, the trust engine 110 serves the requested sensitive data, such as, for example, to the appropriate password to the requesting application. According to another embodiment, the trust engine 110 may include a separate system for storing the sensitive data vault. For example, the trust engine 110 may include a stand-alone software engine implementing the data vault functionality and figuratively residing "behind" the foregoing front-end security system of the trust engine 110. According to this embodiment, the software engine serves the requested sensitive data after the software engine receives a signal indicating positive user authentication from the trust engine 110.

In yet another embodiment, the data vault may be implemented by a third-party system. Similar to the software engine embodiment, the third-party system may advantageously serve the requested sensitive data after the third-party system receives a signal indicating positive user authentication from the trust engine 110. According to yet another embodiment, the data vault may be implemented on the user system 105. A user-side software engine may advantageously serve the foregoing data after receiving a signal indicating positive user authentication from the trust engine 110.

Although the foregoing data vaults are disclosed with reference to alternative embodiments, a skilled artisan will recognize from the disclosure herein, a wide number of additional implementations thereof. For example, a particular data vault may include aspects from some or all of the foregoing embodiments. In addition, any of the foregoing data vaults may employ one or more authentication requests at varying times. For example, any of the data vaults may require authentication every one or more transactions, periodically, every one or more sessions, every access to one or more Webpages or Websites, at one or more other specified intervals, or the like.

FIG. 11 illustrates a data flow of a signing process 1100 according to aspects of an embodiment of the invention. As shown in FIG. 11, the signing process 1100 includes steps similar to those of the authentication process 1000 described in the foregoing with reference to FIG. 10. According to one embodiment of the invention, the signing process 1100 first authenticates the user and then performs one or more of several digital signing functions as will be discussed in further detail below. According to another embodiment, the signing process 1100 may advantageously store data related thereto, such as hashes of messages or documents, or the like. This data may advantageously be used in an audit or any other

event, such as for example, when a participating party attempts to repudiate a transaction.

As shown in FIG. 11, during the authentication steps, the user and vendor may advantageously agree on a message, such as, for example, a contract. During signing, the signing process 1100 advantageously ensures that the contract signed by the user is identical to the contract supplied by the vendor. Therefore, according to one embodiment, during authentication, the vendor and the user include a hash of their respective copies of the message or contract, in the data transmitted to the authentication engine 215. By employing only a hash of a message or contract, the trust engine 110 may advantageously store a significantly reduced amount of data, providing for a more efficient and cost effective cryptographic system. In addition, the stored hash may be advantageously compared to a hash of a document in question to determine whether the document in question matches one signed by any of the parties. The ability to determine whether the document is identical to one relating to a transaction provides for additional evidence that can be used against a claim for repudiation by a party to a transaction.

In step 1103, the authentication engine 215 assembles the enrollment authentication data and compares it to the current authentication data provided by the user. When the comparator of the authentication engine 215 indicates that the enrollment authentication data matches the current authentication data, the comparator of the authentication engine 215 also compares the hash of the message supplied by the vendor to the hash of the message supplied by the user. Thus, the authentication engine 215 advantageously ensures that the message agreed to by the user is identical to that agreed to by the vendor.

In step 1105, the authentication engine 215 transmits a digital signature request to the cryptographic engine 220. According to one embodiment of the invention, the request includes a hash of the message or contract. However, a skilled artisan will recognize from the disclosure herein that the cryptographic engine 220 may encrypt virtually any type of data, including, but not limited to, video, audio, biometrics, images or text to form the desired digital signature. Returning to step 1105, the digital signature request preferably comprises an XML document communicated through conventional SSL technologies.

In step 1110, the authentication engine 215 transmits a request to each of the data storage facilities D1 through D4, such that each of the data storage facilities D1 through D4 transmit their respective portion of the cryptographic key or keys corresponding to a signing party. According to another embodiment, the cryptographic engine 220 employs some or all of the steps of the interoperability process 970 discussed in the foregoing, such that the cryptographic engine 220 first determines the appropriate key or keys to request from the depository 210 or the depository system 700 for the signing party, and takes actions to provide appropriate matching keys. According to still another embodiment, the authentication engine 215 or the cryptographic engine 220 may advantageously request one or more of the keys associated with the signing party and stored in the depository 210 or depository system 700.

According to one embodiment, the signing party includes one or both the user and the vendor. In such case, the authentication engine 215 advantageously requests the cryptographic keys corresponding to the user and/or the vendor. According to another embodiment, the signing party includes the trust engine 110. In this embodiment, the trust engine 110 is certifying that the authentication process 1000 properly authenticated the user, vendor, or both. Therefore, the authentication engine 215

requests the cryptographic key of the trust engine 110, such as, for example, the key belonging to the cryptographic engine 220, to perform the digital signature. According to another embodiment, the trust engine 110 performs a digital notary-like function. In this embodiment, the signing party includes the user, vendor, or both, along with the trust engine 110. Thus, the trust engine 110 provides the digital signature of the user and/or vendor, and then indicates with its own digital signature that the user and/or vendor were properly authenticated. In this embodiment, the authentication engine 215 may advantageously request assembly of the cryptographic keys corresponding to the user, the vendor, or both. According to another embodiment, the authentication engine 215 may advantageously request assembly of the cryptographic keys corresponding to the trust engine 110.

According to another embodiment, the trust engine 110 performs power of attorney-like functions. For example, the trust engine 110 may digitally sign the message on behalf of a third party. In such case, the authentication engine 215 requests the cryptographic keys associated with the third party. According to this embodiment, the signing process 1100 may advantageously include authentication of the third party, before allowing power of attorney-like functions. In addition, the authentication process 1000 may include a check for third party constraints, such as, for example, business logic or the like dictating when and in what circumstances a particular third-party's signature may be used.

Based on the foregoing, in step 1110, the authentication engine requested the cryptographic keys from the data storage facilities D1 through D4 corresponding to the signing party. In step 1115, the data storage facilities D1 through D4 transmit their respective portions of the cryptographic key corresponding to the signing party to the cryptographic engine 220. According to one embodiment, the foregoing transmissions include SSL technologies. According to another embodiment, the foregoing transmissions may advantageously be super-encrypted with the public key of the cryptographic engine 220.

In step 1120, the cryptographic engine 220 assembles the foregoing cryptographic keys of the signing party and encrypts the message therewith, thereby forming the digital signature(s). In step 1125 of the signing process 1100, the cryptographic engine 220 transmits the digital signature(s) to the authentication engine 215. In step 1130, the authentication engine 215 transmits the filled-in authentication request along with a copy of the hashed message and the digital signature(s) to the transaction engine 205. In step 1135, the transaction engine 205 transmits a receipt comprising the transaction ID, an indication of whether the authentication was successful, and the digital signature(s), to the vendor. According to one embodiment, the foregoing transmission may advantageously include the digital signature of the trust engine 110. For example, the trust engine 110 may encrypt the hash of the receipt with its private key, thereby forming a digital signature to be attached to the transmission to the vendor.

According to one embodiment, the transaction engine 205 also transmits a confirmation message to the user. Although the signing process 1100 is disclosed with reference to its preferred and alternative embodiments, the invention is not intended to be limited thereby. Rather, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for the signing process 1100. For example, the vendor may be replaced with a user application, such as an email application. For example, the user may wish to digitally sign a particular email with his or her digital signature. In such an embodiment, the transmission throughout the signing process

1100 may advantageously include only one copy of a hash of the message. Moreover, a skilled artisan will recognize from the disclosure herein that a wide number of client applications may request digital signatures. For example, the client applications may comprise word processors, spreadsheets, emails, voicemail, access to restricted system areas, or the like.

In addition, a skilled artisan will recognize from the disclosure herein that steps **1105** through **1120** of the signing process **1100** may advantageously employ some or all of the steps of the interoperability process **970** of FIG. 9B, thereby providing interoperability between differing cryptographic systems that may, for example, need to process the digital signature under differing signature types.

FIG. 12 illustrates a data flow of an encryption/decryption process **1200** according to aspects of an embodiment of the invention. As shown in FIG. 12, the decryption process **1200** begins by authenticating the user using the authentication process **1000**. According to one embodiment, the authentication process **1000** includes in the authentication request, a synchronous session key. For example, in conventional PKI technologies, it is understood by skilled artisans that encrypting or decrypting data using public and private keys is mathematically intensive and may require significant system resources. However, in symmetric key cryptographic systems, or systems where the sender and receiver of a message share a single common key that is used to encrypt and decrypt a message, the mathematical operations are significantly simpler and faster. Thus, in the conventional PKI technologies, the sender of a message will generate synchronous session key, and encrypt the message using the simpler, faster symmetric key system. Then, the sender will encrypt the session key with the public key of the receiver. The encrypted session key will be attached to the synchronously encrypted message and both data are sent to the receiver. The receiver uses his or her private key to decrypt the session key, and then uses the session key to decrypt the message. Based on the foregoing, the simpler and faster symmetric key system is used for the majority of the encryption/decryption processing. Thus, in the decryption process **1200**, the decryption advantageously assumes that a synchronous key has been encrypted with the public key of the user. Thus, as mentioned in the foregoing, the encrypted session key is included in the authentication request.

Returning to the decryption process **1200**, after the user has been authenticated in step **1205**, the authentication engine **215** forwards the encrypted session key to the cryptographic engine **220**. In step **1210**, the authentication engine **215** forwards a request to each of the data storage facilities, D1 through D4, requesting the cryptographic key data of the user. In step **1215**, each data storage facility, D1 through D4, transmits their respective portion of the cryptographic key to the cryptographic engine **220**. According to one embodiment, the foregoing transmission is encrypted with the public key of the cryptographic engine **220**.

In step **1220** of the decryption process **1200**, the cryptographic engine **220** assembles the cryptographic key and decrypts the session key therewith. In step **1225**, the cryptographic engine forwards the session key to the authentication engine **215**. In step **1227**, the authentication engine **215** fills in the authentication request including the decrypted session key, and transmits the filled-in authentication request to the transaction engine **205**. In step **1230**, the transaction engine **205** forwards the authentication request along with the session key to the requesting application or vendor. Then, according to one embodiment, the requesting application or vendor uses the session key to decrypt the encrypted message.

Although the decryption process **1200** is disclosed with reference to its preferred and alternative embodiments, a skilled artisan will recognize from the disclosure herein, a wide number of alternatives for the decryption process **1200**. For example, the decryption process **1200** may forego synchronous key encryption and rely on full public-key technology. In such an embodiment, the requesting application may transmit the entire message to the cryptographic engine **220**, or, may employ some type of compression or reversible hash in order to transmit the message to the cryptographic engine **220**. A skilled artisan will also recognize from the disclosure herein that the foregoing communications may advantageously include XML documents wrapped in SSL technology.

The encryption/decryption process **1200** also provides for encryption of documents or other data. Thus, in step **1235**, a requesting application or vendor may advantageously transmit to the transaction engine **205** of the trust engine **110**, a request for the public key of the user. The requesting application or vendor makes this request because the requesting application or vendor uses the public key of the user, for example, to encrypt the session key that will be used to encrypt the document or message. As mentioned in the enrollment process **900**, the transaction engine **205** stores a copy of the digital certificate of the user, for example, in the mass storage **225**. Thus, in step **1240** of the encryption process **1200**, the transaction engine **205** requests the digital certificate of the user from the mass storage **225**. In step **1245**, the mass storage **225** transmits the digital certificate corresponding to the user, to the transaction engine **205**. In step **1250**, the transaction engine **205** transmits the digital certificate to the requesting application or vendor. According to one embodiment, the encryption portion of the encryption process **1200** does not include the authentication of a user. This is because the requesting vendor needs only the public key of the user, and is not requesting any sensitive data.

A skilled artisan will recognize from the disclosure herein that if a particular user does not have a digital certificate, the trust engine **110** may employ some or all of the enrollment process **900** in order to generate a digital certificate for that particular user. Then, the trust engine **110** may initiate the encryption/decryption process **1200** and thereby provide the appropriate digital certificate. In addition, a skilled artisan will recognize from the disclosure herein that steps **1220** and **1235** through **1250** of the encryption/decryption process **1200** may advantageously employ some or all of the steps of the interoperability process of FIG. 9B, thereby providing interoperability between differing cryptographic systems that may, for example, need to process the encryption.

FIG. 13 illustrates a simplified block diagram of a trust engine system **1300** according to aspects of yet another embodiment of the invention. As shown in FIG. 13, the trust engine system **1300** comprises a plurality of distinct trust engines **1305**, **1310**, **1315**, and **1320**, respectively. To facilitate a more complete understanding of the invention, FIG. 13 illustrates each trust engine, **1305**, **1310**, **1315**, and **1320** as having a transaction engine, a depository, and an authentication engine. However, a skilled artisan will recognize that each transaction engine may advantageously comprise some, a combination, or all of the elements and communication channels disclosed with reference to FIGS. 1-8. For example, one embodiment may advantageously include trust engines having one or more transaction engines, depositories, and cryptographic servers or any combinations thereof.

According to one embodiment of the invention, each of the trust engines **1305**, **1310**, **1315** and **1320** are geographically separated, such that, for example, the trust engine **1305** may

31

reside in a first location, the trust engine **1310** may reside in a second location, the trust engine **1315** may reside in a third location, and the trust engine **1320** may reside in a fourth location. The foregoing geographic separation advantageously decreases system response time while increasing the security of the overall trust engine system **1300**.

For example, when a user logs onto the cryptographic system **100**, the user may be nearest the first location and may desire to be authenticated. As described with reference to FIG. **10**, to be authenticated, the user provides current authentication data, such as a biometric or the like, and the current authentication data is compared to that user's enrollment authentication data. Therefore, according to one example, the user advantageously provides current authentication data to the geographically nearest trust engine **1305**. The transaction engine **1321** of the trust engine **1305** then forwards the current authentication data to the authentication engine **1322** also residing at the first location. According to another embodiment, the transaction engine **1321** forwards the current authentication data to one or more of the authentication engines of the trust engines **1310**, **1315**, or **1320**.

The transaction engine **1321** also requests the assembly of the enrollment authentication data from the depositories of, for example, each of the trust engines, **1305** through **1320**. According to this embodiment, each depository provides its portion of the enrollment authentication data to the authentication engine **1322** of the trust engine **1305**. The authentication engine **1322** then employs the encrypted data portions from, for example, the first two depositories to respond, and assembles the enrollment authentication data into deciphered form. The authentication engine **1322** compares the enrollment authentication data with the current authentication data and returns an authentication result to the transaction engine **1321** of the trust engine **1305**.

Based on the above, the trust engine system **1300** employs the nearest one of a plurality of geographically separated trust engines, **1305** through **1320**, to perform the authentication process. According to one embodiment of the invention, the routing of information to the nearest transaction engine may advantageously be performed at client-side applets executing on one or more of the user system **105**, vendor system **120**, or certificate authority **115**. According to an alternative embodiment, a more sophisticated decision process may be employed to select from the trust engines **1305** through **1320**. For example, the decision may be based on the availability, operability, speed of connections, load, performance, geographic proximity, or a combination thereof, of a given trust engine.

In this way, the trust engine system **1300** lowers its response time while maintaining the security advantages associated with geographically remote data storage facilities, such as those discussed with reference to FIG. **7** where each data storage facility stores randomized portions of sensitive data. For example, a security compromise at, for example, the depository **1325** of the trust engine **1315** does not necessarily compromise the sensitive data of the trust engine system **1300**. This is because the depository **1325** contains only non-decipherable randomized data that, without more, is entirely useless.

According to another embodiment, the trust engine system **1300** may advantageously include multiple cryptographic engines arranged similar to the authentication engines. The cryptographic engines may advantageously perform cryptographic functions such as those disclosed with reference to FIGS. **1-8**. According to yet another embodiment, the trust engine system **1300** may advantageously replace the multiple authentication engines with multiple cryptographic engines,

32

thereby performing cryptographic functions such as those disclosed with reference to FIGS. **1-8**. According to yet another embodiment of the invention, the trust engine system **1300** may replace each multiple authentication engine with an engine having some or all of the functionality of the authentication engines, cryptographic engines, or both, as disclosed in the foregoing.

Although the trust engine system **1300** is disclosed with reference to its preferred and alternative embodiments, a skilled artisan will recognize that the trust engine system **1300** may comprise portions of trust engines **1305** through **1320**. For example, the trust engine system **1300** may include one or more transaction engines, one or more depositories, one or more authentication engines, or one or more cryptographic engines or combinations thereof.

FIG. **14** illustrates a simplified block diagram of a trust engine System **1400** according to aspects of yet another embodiment of the invention. As shown in FIG. **14**, the trust engine system **1400** includes multiple trust engines **1405**, **1410**, **1415** and **1420**. According to one embodiment, each of the trust engines **1405**, **1410**, **1415** and **1420**, comprise some or all of the elements of trust engine **110** disclosed with reference to FIGS. **1-8**. According to this embodiment, when the client side applets of the user system **105**, the vendor system **120**, or the certificate authority **115**, communicate with the trust engine system **1400**, those communications are sent to the IP address of each of the trust engines **1405** through **1420**. Further, each transaction engine of each of the trust engines, **1405**, **1410**, **1415**, and **1420**, behaves similar to the transaction engine **1321** of the trust engine **1305** disclosed with reference to FIG. **13**. For example, during an authentication process, each transaction engine of each of the trust engines **1405**, **1410**, **1415**, and **1420** transmits the current authentication data to their respective authentication engines and transmits a request to assemble the randomized data stored in each of the depositories of each of the trust engines **1405** through **1420**. FIG. **14** does not illustrate all of these communications; as such illustration would become overly complex. Continuing with the authentication process, each of the depositories then communicates its portion of the randomized data to each of the authentication engines of the each of the trust engines **1405** through **1420**. Each of the authentication engines of the each of the trust engines employs its comparator to determine whether the current authentication data matches the enrollment authentication data provided by the depositories of each of the trust engines **1405** through **1420**. According to this embodiment, the result of the comparison by each of the authentication engines is then transmitted to a redundancy module of the other three trust engines. For example, the result of the authentication engine from the trust engine **1405** is transmitted to the redundancy modules of the trust engines **1410**, **1415**, and **1420**. Thus, the redundancy module of the trust engine **1405** likewise receives the result of the authentication engines from the trust engines **1410**, **1415**, and **1420**.

FIG. **15** illustrates a block diagram of the redundancy module of FIG. **14**. The redundancy module comprises a comparator configured to receive the authentication result from three authentication engines and transmit that result to the transaction engine of the fourth trust engine. The comparator compares the authentication result from the three authentication engines, and if two of the results agree, the comparator concludes that the authentication result should match that of the two agreeing authentication engines. This result is then transmitted back to the transaction engine corresponding to the trust engine not associated with the three authentication engines.

Based on the foregoing, the redundancy module determines an authentication result from data received from authentication engines that are preferably geographically remote from the trust engine of that the redundancy module. By providing such redundancy functionality, the trust engine system **1400** ensures that a compromise of the authentication engine of one of the trust engines **1405** through **1420**, is insufficient to compromise the authentication result of the redundancy module of that particular trust engine. A skilled artisan will recognize that redundancy module functionality of the trust engine system **1400** may also be applied to the cryptographic engine of each of the trust engines **1405** through **1420**. However, such cryptographic engine communication was not shown in FIG. **14** to avoid complexity. Moreover, a skilled artisan will recognize a wide number of alternative authentication result conflict resolution algorithms for the comparator of FIG. **15** are suitable for use in the present invention.

According to yet another embodiment of the invention, the trust engine system **1400** may advantageously employ the redundancy module during cryptographic comparison steps. For example, some or all of the foregoing redundancy module disclosure with reference to FIGS. **14** and **15** may advantageously be implemented during a hash comparison of documents provided by one or more parties during a particular transaction.

Although the foregoing invention has been described in terms of certain preferred and alternative embodiments, other embodiments will be apparent to those of ordinary skill in the art from the disclosure herein. For example, the trust engine **110** may issue short-term certificates, where the private cryptographic key is released to the user for a predetermined period of time. For example, current certificate standards include a validity field that can be set to expire after a predetermined amount of time. Thus, the trust engine **110** may release a private key to a user where the private key would be valid for, for example, 24 hours. According to such an embodiment, the trust engine **110** may advantageously issue a new cryptographic key pair to be associated with a particular user and then release the private key of the new cryptographic key pair. Then, once the private cryptographic key is released, the trust engine **110** immediately expires any internal valid use of such private key, as it is no longer securable by the trust engine **110**.

In addition, a skilled artisan will recognize that the cryptographic system **100** or the trust engine **110** may include the ability to recognize any type of devices, such as, but not limited to, a laptop, a cell phone, a network, a biometric device or the like. According to one embodiment, such recognition may come from data supplied in the request for a particular service, such as, a request for authentication leading to access or use, a request for cryptographic functionality, or the like. According to one embodiment, the foregoing request may include a unique device identifier, such as, for example, a processor ID. Alternatively, the request may include data in a particular recognizable data format. For example, mobile and satellite phones often do not include the processing power for full X509.v3 heavy encryption certificates, and therefore do not request them. According to this embodiment, the trust engine **110** may recognize the type of data format presented, and respond only in kind.

In an additional aspect of the system described above, context sensitive authentication can be provided using various techniques as will be described below. Context sensitive authentication, for example as shown in FIG. **16**, provides the possibility of evaluating not only the actual data which is sent by the user when attempting to authenticate himself, but also

the circumstances surrounding the generation and delivery of that data. Such techniques may also support transaction specific trust arbitrage between the user and trust engine **110** or between the vendor and trust engine **110**, as will be described below.

As discussed above, authentication is the process of proving that a user is who he says he is. Generally, authentication requires demonstrating some fact to an authentication authority. The trust engine **110** of the present invention represents the authority to which a user must authenticate himself. The user must demonstrate to the trust engine **110** that he is who he says he is by either: knowing something that only the user should know (knowledge-based authentication), having something that only the user should have (token-based authentication), or by being something that only the user should be (biometric-based authentication).

Examples of knowledge-based authentication include without limitation a password, PIN number, or lock combination. Examples of token-based authentication include without limitation a house key, a physical credit card, a driver's license, or a particular phone number. Examples of biometric-based authentication include without limitation a fingerprint, handwriting analysis, facial scan, hand scan, ear scan, iris scan, vascular pattern, DNA, a voice analysis, or a retinal scan.

Each type of authentication has particular advantages and disadvantages, and each provides a different level of security. For example, it is generally harder to create a false fingerprint that matches someone else's than it is to overhear someone's password and repeat it. Each type of authentication also requires a different type of data to be known to the authenticating authority in order to verify someone using that form of authentication.

As used herein, "authentication" will refer broadly to the overall process of verifying someone's identity to be who he says he is. An "authentication technique" will refer to a particular type of authentication based upon a particular piece of knowledge, physical token, or biometric reading. "Authentication data" refers to information which is sent to or otherwise demonstrated to an authentication authority in order to establish identity. "Enrollment data" will refer to the data which is initially submitted to an authentication authority in order to establish a baseline for comparison with authentication data. An "authentication instance" will refer to the data associated with an attempt to authenticate by an authentication technique.

The internal protocols and communications involved in the process of authenticating a user is described with reference to FIG. **10** above. The part of this process within which the context sensitive authentication takes place occurs within the comparison step shown as step **1045** of FIG. **10**. This step takes place within the authentication engine **215** and involves assembling the enrollment data **410** retrieved from the depository **210** and comparing the authentication data provided by the user to it. One particular embodiment of this process is shown in FIG. **16** and described below.

The current authentication data provided by the user and the enrollment data retrieved from the depository **210** are received by the authentication engine **215** in step **1600** of FIG. **16**. Both of these sets of data may contain data which is related to separate techniques of authentication. The authentication engine **215** separates the authentication data associated with each individual authentication instance in step **1605**. This is necessary so that the authentication data is compared with the appropriate subset of the enrollment data

for the user (e.g. fingerprint authentication data should be compared with fingerprint enrollment data, rather than password enrollment data).

Generally, authenticating a user involves one or more individual authentication instances, depending on which authentication techniques are available to the user. These methods are limited by the enrollment data which were provided by the user during his enrollment process (if the user did not provide a retinal scan when enrolling, he will not be able to authenticate himself using a retinal scan), as well as the means which may be currently available to the user (e.g. if the user does not have a fingerprint reader at his current location, fingerprint authentication will not be practical). In some cases, a single authentication instance may be sufficient to authenticate a user; however, in certain circumstances a combination of multiple authentication instances may be used in order to more confidently authenticate a user for a particular transaction.

Each authentication instance consists of data related to a particular authentication technique (e.g. fingerprint, password, smart card, etc.) and the circumstances which surround the capture and delivery of the data for that particular technique. For example, a particular instance of attempting to authenticate via password will generate not only the data related to the password itself, but also circumstantial data, known as "metadata", related to that password attempt. This circumstantial data includes information such as: the time at which the particular authentication instance took place, the network address from which the authentication information was delivered, as well as any other information as is known to those of skill in the art which may be determined about the origin of the authentication data (the type of connection, the processor serial number, etc.).

In many cases, only a small amount of circumstantial metadata will be available. For example, if the user is located on a network which uses proxies or network address translation or another technique which masks the address of the originating computer, only the address of the proxy or router may be determined. Similarly, in many cases information such as the processor serial number will not be available because of either limitations of the hardware or operating system being used, disabling of such features by the operator of the system, or other limitations of the connection between the user's system and the trust engine **110**.

As shown in FIG. **16**, once the individual authentication instances represented within the authentication data are extracted and separated in step **1605**, the authentication engine **215** evaluates each instance for its reliability in indicating that the user is who he claims to be. The reliability for a single authentication instance will generally be determined based on several factors. These may be grouped as factors relating to the reliability associated with the authentication technique, which are evaluated in step **1610**, and factors relating to the reliability of the particular authentication data provided, which are evaluated in step **1815**. The first group includes without limitation the inherent reliability of the authentication technique being used, and the reliability of the enrollment data being used with that method. The second group includes without limitation the degree of match between the enrollment data and the data provided with the authentication instance, and the metadata associated with that authentication instance. Each of these factors may vary independently of the others.

The inherent reliability of an authentication technique is based on how hard it is for an imposter to provide someone else's correct data, as well as the overall error rates for the authentication technique. For passwords and knowledge

based authentication methods, this reliability is often fairly low because there is nothing that prevents someone from revealing their password to another person and for that second person to use that password. Even a more complex knowledge based system may have only moderate reliability since knowledge may be transferred from person to person fairly easily. Token based authentication, such as having a proper smart card or using a particular terminal to perform the authentication, is similarly of low reliability used by itself, since there is no guarantee that the right person is in possession of the proper token.

However, biometric techniques are more inherently reliable because it is generally difficult to provide someone else with the ability to use your fingerprints in a convenient manner, even intentionally. Because subverting biometric authentication techniques is more difficult, the inherent reliability of biometric methods is generally higher than that of purely knowledge or token based authentication techniques. However, even biometric techniques may have some occasions in which a false acceptance or false rejection is generated. These occurrences may be reflected by differing reliabilities for different implementations of the same biometric technique. For example, a fingerprint matching system provided by one company may provide a higher reliability than one provided by a different company because one uses higher quality optics or a better scanning resolution or some other improvement which reduces the occurrence of false acceptances or false rejections.

Note that this reliability may be expressed in different manners. The reliability is desirably expressed in some metric which can be used by the heuristics **530** and algorithms of the authentication engine **215** to calculate the confidence level of each authentication. One preferred mode of expressing these reliabilities is as a percentage or fraction. For instance, fingerprints might be assigned an inherent reliability of 97%, while passwords might only be assigned an inherent reliability of 50%. Those of skill in the art will recognize that these particular values are merely exemplary and may vary between specific implementations.

The second factor for which reliability must be assessed is the reliability of the enrollment. This is part of the "graded enrollment" process referred to above. This reliability factor reflects the reliability of the identification provided during the initial enrollment process. For instance, if the individual initially enrolls in a manner where they physically produce evidence of their identity to a notary or other public official, and enrollment data is recorded at that time and notarized, the data will be more reliable than data which is provided over a network during enrollment and only vouched for by a digital signature or other information which is not truly tied to the individual.

Other enrollment techniques with varying levels of reliability include without limitation: enrollment at a physical office of the trust engine **110** operator; enrollment at a user's place of employment; enrollment at a post office or passport office; enrollment through an affiliated or trusted party to the trust engine **110** operator; anonymous or pseudonymous enrollment in which the enrolled identity is not yet identified with a particular real individual, as well as such other means as are known in the art.

These factors reflect the trust between the trust engine **110** and the source of identification provided during the enrollment process. For instance, if enrollment is performed in association with an employer during the initial process of providing evidence of identity, this information may be considered extremely reliable for purposes within the company, but may be trusted to a lesser degree by a government agency,

or by a competitor. Therefore, trust engines operated by each of these other organizations may assign different levels of reliability to this enrollment.

Similarly, additional data which is submitted across a network, but which is authenticated by other trusted data provided during a previous enrollment with the same trust engine **110** may be considered as reliable as the original enrollment data was, even though the latter data were submitted across an open network. In such circumstances, a subsequent notarization will effectively increase the level of reliability associated with the original enrollment data. In this way for example, an anonymous or pseudonymous enrollment may then be raised to a full enrollment by demonstrating to some enrollment official the identity of the individual matching the enrolled data.

The reliability factors discussed above are generally values which may be determined in advance of any particular authentication instance. This is because they are based upon the enrollment and the technique, rather than the actual authentication. In one embodiment, the step of generating reliability based upon these factors involves looking up previously determined values for this particular authentication technique and the enrollment data of the user. In a further aspect of an advantageous embodiment of the present invention, such reliabilities may be included with the enrollment data itself. In this way, these factors are automatically delivered to the authentication engine **215** along with the enrollment data sent from the depository **210**.

While these factors may generally be determined in advance of any individual authentication instance, they still have an effect on each authentication instance which uses that particular technique of authentication for that user. Furthermore, although the values may change over time (e.g. if the user re-enrolls in a more reliable fashion), they are not dependent on the authentication data itself. By contrast, the reliability factors associated with a single specific instance's data may vary on each occasion. These factors, as discussed below, must be evaluated for each new authentication in order to generate reliability scores in step **1815**.

The reliability of the authentication data reflects the match between the data provided by the user in a particular authentication instance and the data provided during the authentication enrollment. This is the fundamental question of whether the authentication data matches the enrollment data for the individual the user is claiming to be. Normally, when the data do not match, the user is considered to not be successfully authenticated, and the authentication fails. The manner in which this is evaluated may change depending on the authentication technique used. The comparison of such data is performed by the comparator **515** function of the authentication engine **215** as shown in FIG. **5**.

For instance, matches of passwords are generally evaluated in a binary fashion. In other words, a password is either a perfect match, or a failed match. It is usually not desirable to accept as even a partial match a password which is close to the correct password if it is not exactly correct. Therefore, when evaluating a password authentication, the reliability of the authentication returned by the comparator **515** is typically either 100% (correct) or 0% (wrong), with no possibility of intermediate values.

Similar rules to those for passwords are generally applied to token based authentication methods, such as smart cards. This is because having a smart card which has a similar identifier or which is similar to the correct one, is still just as wrong as having any other incorrect token. Therefore tokens tend also to be binary authenticators: a user either has the right token, or he doesn't.

However, certain types of authentication data, such as questionnaires and biometrics, are generally not binary authenticators. For example, a fingerprint may match a reference fingerprint to varying degrees. To some extent, this may be due to variations in the quality of the data captured either during the initial enrollment or in subsequent authentications. (A fingerprint may be smudged or a person may have a still healing scar or burn on a particular finger.) In other instances the data may match less than perfectly because the information itself is somewhat variable and based upon pattern matching. (A voice analysis may seem close but not quite right because of background noise, or the acoustics of the environment in which the voice is recorded, or because the person has a cold.) Finally, in situations where large amounts of data are being compared, it may simply be the case that much of the data matches well, but some doesn't. (A ten-question questionnaire may have resulted in eight correct answers to personal questions, but two incorrect answers.) For any of these reasons, the match between the enrollment data and the data for a particular authentication instance may be desirably assigned a partial match value by the comparator **515**. In this way, the fingerprint might be said to be a 85% match, the voice print a 65% match, and the questionnaire an 80% match, for example.

This measure (degree of match) produced by the comparator **515** is the factor representing the basic issue of whether an authentication is correct or not. However, as discussed above, this is only one of the factors which may be used in determining the reliability of a given authentication instance. Note also that even though a match to some partial degree may be determined, that ultimately, it may be desirable to provide a binary result based upon a partial match. In an alternate mode of operation, it is also possible to treat partial matches as binary, i.e. either perfect (100%) or failed (0%) matches, based upon whether or not the degree of match passes a particular threshold level of match. Such a process may be used to provide a simple pass/fail level of matching for systems which would otherwise produce partial matches.

Another factor to be considered in evaluating the reliability of a given authentication instance concerns the circumstances under which the authentication data for this particular instance are provided. As discussed above, the circumstances refer to the metadata associated with a particular authentication instance. This may include without limitation such information as: the network address of the authenticator, to the extent that it can be determined; the time of the authentication; the mode of transmission of the authentication data (phone line, cellular, network, etc.); and the serial number of the system of the authenticator.

These factors can be used to produce a profile of the type of authentication that is normally requested by the user. Then, this information can be used to assess reliability in at least two manners. One manner is to consider whether the user is requesting authentication in a manner which is consistent with the normal profile of authentication by this user. If the user normally makes authentication requests from one network address during business days (when she is at work) and from a different network address during evenings or weekends (when she is at home), an authentication which occurs from the home address during the business day is less reliable because it is outside the normal authentication profile. Similarly, if the user normally authenticates using a fingerprint biometric and in the evenings, an authentication which originates during the day using only a password is less reliable.

An additional way in which the circumstantial metadata can be used to evaluate the reliability of an instance of authentication is to determine how much corroboration the circum-

stance provides that the authenticator is the individual he claims to be. For instance, if the authentication comes from a system with a serial number known to be associated with the user, this is a good circumstantial indicator that the user is who they claim to be. Conversely, if the authentication is coming from a network address which is known to be in Los Angeles when the user is known to reside in London, this is an indication that this authentication is less reliable based on its circumstances.

It is also possible that a cookie or other electronic data may be placed upon the system being used by a user when they interact with a vendor system or with the trust engine **110**. This data is written to the storage of the system of the user and may contain an identification which may be read by a Web browser or other software on the user system. If this data is allowed to reside on the user system between sessions (a "persistent cookie"), it may be sent with the authentication data as further evidence of the past use of this system during authentication of a particular user. In effect, the metadata of a given instance, particularly a persistent cookie, may form a sort of token based authenticator itself.

Once the appropriate reliability factors based on the technique and data of the authentication instance are generated as described above in steps **1610** and **1615** respectively, they are used to produce an overall reliability for the authentication instance provided in step **1620**. One means of doing this is simply to express each reliability as a percentage and then to multiply them together.

For example, suppose the authentication data is being sent in from a network address known to be the user's home computer completely in accordance with the user's past authentication profile (100%), and the technique being used is fingerprint identification (97%), and the initial finger print data was roistered through the user's employer with the trust engine **110** (90%), and the match between the authentication data and the original fingerprint template in the enrollment data is very good (99%). The overall reliability of this authentication instance could then be calculated as the product of these reliabilities: $100\% * 97\% * 90\% * 99\% = 86.4\%$ reliability.

This calculated reliability represents the reliability of one single instance of authentication. The overall reliability of a single authentication instance may also be calculated using techniques which treat the different reliability factors differently, for example by using formulas where different weights are assigned to each reliability factor. Furthermore, those of skill in the art will recognize that the actual values used may represent values other than percentages and may use non-arithmetic systems. One embodiment may include a module used by an authentication requestor to set the weights for each factor and the algorithms used in establishing the overall reliability of the authentication instance.

The authentication engine **215** may use the above techniques and variations thereof to determine the reliability of a single authentication instance, indicated as step **1620**. However, it may be useful in many authentication situations for multiple authentication instances to be provided at the same time. For example, while attempting to authenticate himself using the system of the present invention, a user may provide a user identification, fingerprint authentication data, a smart card, and a password. In such a case, three independent authentication instances are being provided to the trust engine **110** for evaluation. Proceeding to step **1625**, if the authentication engine **215** determines that the data provided by the user includes more than one authentication instance, then each instance in turn will be selected as shown in step **1630** and evaluated as described above in steps **1610**, **1615** and **1620**.

Note that many of the reliability factors discussed may vary from one of these instances to another. For instance, the inherent reliability of these techniques is likely to be different, as well as the degree of match provided between the authentication data and the enrollment data. Furthermore, the user may have provided enrollment data at different times and under different circumstances for each of these techniques, providing different enrollment reliabilities for each of these instances as well. Finally, even though the circumstances under which the data for each of these instances is being submitted is the same, the use of such techniques may each fit the profile of the user differently, and so may be assigned different circumstantial reliabilities. (For example, the user may normally use their password and fingerprint, but not their smart card.)

As a result, the final reliability for each of these authentication instances may be different from One another. However, by using multiple instances together, the overall confidence level for the authentication will tend to increase.

Once the authentication engine has performed steps **1610** through **1620** for all of the authentication instances provided in the authentication data, the reliability of each instance is used in step **1635** to evaluate the overall authentication confidence level. This process of combining the individual authentication instance reliabilities into the authentication confidence level may be modeled by various methods relating the individual reliabilities produced, and may also address the particular interaction between some of these authentication techniques. (For example, multiple knowledge-based systems such as passwords may produce less confidence than a single password and even a fairly weak biometric, such as a basic voice analysis.)

One means in which the authentication engine **215** may combine the reliabilities of multiple concurrent authentication instances to generate a final confidence level is to multiply the unreliability of each instance to arrive at a total unreliability. The unreliability is generally the complementary percentage of the reliability. For example, a technique which is 84% reliable is 16% unreliable. The three authentication instances described above (fingerprint, smart card, password) which produce reliabilities of 86%, 75%, and 72% would have corresponding unreliabilities of $(100-86)\%$, $(100-75)\%$ and $(100-72)\%$, or 14%, 25%, and 28%, respectively. By multiplying these unreliabilities, we get a cumulative unreliability of $14\% * 25\% * 28\% = 0.98\%$ unreliability, which corresponds to a reliability of 99.02%.

In an additional mode of operation, additional factors and heuristics **530** may be applied within the authentication engine **215** to account for the interdependence of various authentication techniques. For example, if someone has unauthorized access to a particular home computer, they probably have access to the phone line at that address as well. Therefore, authenticating based on an originating phone number as well as upon the serial number of the authenticating system does not add much to the overall confidence in the authentication. However, knowledge based authentication is largely independent of token based authentication (i.e. if someone steals your cellular phone or keys, they are no more likely to know your PIN or password than if they hadn't).

Furthermore, different vendors or other authentication requestors may wish to weigh different aspects of the authentication differently. This may include the use of separate weighing factors or algorithms used in calculating the reliability of individual instances as well as the use of different means to evaluate authentication events with multiple instances.

For instance, vendors for certain types of transactions, for instance corporate email systems, may desire to authenticate primarily based upon heuristics and other circumstantial data by default. Therefore, they may apply high weights to factors related to the metadata and other profile related information associated with the circumstances surrounding authentication events. This arrangement could be used to ease the burden on users during normal operating hours, by not requiring more from the user than that he be logged on to the correct machine during business hours. However, another vendor may weigh authentications coming from a particular technique most heavily, for instance fingerprint matching, because of a policy decision that such a technique is most suited to authentication for the particular vendor's purposes.

Such varying weights may be defined by the authentication requestor in generating the authentication request and sent to the trust engine 110 with the authentication request in one mode of operation. Such options could also be set as preferences during an initial enrollment process for the authentication requestor and stored within the authentication engine in another mode of operation.

Once the authentication engine 215 produces an authentication confidence level for the authentication data provided, this confidence level is used to complete the authentication request in step 1640, and this information is forwarded from the authentication engine 215 to the transaction engine 205 for inclusion in a message to the authentication requestor.

The process described above is merely exemplary, and those of skill in the art will recognize that the steps need not be performed in the order shown or that only certain of the steps are desired to be performed, or that a variety of combinations of steps may be desired. Furthermore, certain steps, such as the evaluation of the reliability of each authentication instance provided, may be carried out in parallel with one another if circumstances permit.

In a further aspect of this invention, a method is provided to accommodate conditions when the authentication confidence level produced by the process described above fails to meet the required trust level of the vendor or other party requiring the authentication. In circumstances such as these where a gap exists between the level of confidence provided and the level of trust desired, the operator of the trust engine 110 is in a position to provide opportunities for one or both parties to provide alternate data or requirements in order to close this trust gap. This process will be referred to as "trust arbitrage" herein.

Trust arbitrage may take place within a framework of cryptographic authentication as described above with reference to FIGS. 10 and 11. As shown therein, a vendor or other party will request authentication of a particular user in association with a particular transaction. In one circumstance, the vendor simply requests an authentication, either positive or negative, and after receiving appropriate data from the user, the trust engine 110 will provide such a binary authentication. In circumstances such as these, the degree of confidence required in order to secure a positive authentication is determined based upon preferences set within the trust engine 110.

However, it is also possible that the vendor may request a particular level of trust in order to complete a particular transaction. This required level may be included with the authentication request (e.g. authenticate this user to 98% confidence) or may be determined by the trust engine 110 based on other factors associated with the transaction (i.e. authenticate this user as appropriate for this transaction). One such factor might be the economic value of the transaction. For transactions which have greater economic value, a higher degree of trust may be required. Similarly, for transactions with high

degrees of risk a high degree of trust may be required. Conversely, for transactions which are either of low risk or of low value, lower trust levels may be required by the vendor or other authentication requestor.

The process of trust arbitrage occurs between the steps of the trust engine 110 receiving the authentication data in step 1050 of FIG. 10 and the return of an authentication result to the vendor in step 1055 of FIG. 10. Between these steps, the process which leads to the evaluation of trust levels and the potential trust arbitrage occurs as shown in FIG. 17. In circumstances where simple binary authentication is performed, the process shown in FIG. 17 reduces to having the transaction engine 205 directly compare the authentication data provided with the enrollment data for the identified user as discussed above with reference to FIG. 10, flagging any difference as a negative authentication.

As shown in FIG. 17, the first step after receiving the data in step 1050 is for the transaction engine 205 to determine the trust level which is required for a positive authentication for this particular transaction in step 1710. This step may be performed by one of several different methods. The required trust level may be specified to the trust engine 110 by the authentication requestor at the time when the authentication request is made. The authentication requestor may also set a preference in advance which is stored within the depository 210 or other storage which is accessible by the transaction engine 205. This preference may then be read and used each time an authentication request is made by this authentication requestor. The preference may also be associated with a particular user as a security measure such that a particular level of trust is always required in order to authenticate that user, the user preference being stored in the depository 210 or other storage media accessible by the transaction engine 205. The required level may also be derived by the transaction engine 205 or authentication engine 215 based upon information provided in the authentication request, such as the value and risk level of the transaction to be authenticated.

In one mode of operation, a policy management module or other software which is used when generating the authentication request is used to specify the required degree of trust for the authentication of the transaction. This may be used to provide a series of rules to follow when assigning the required level of trust based upon the policies which are specified within the policy management module. One advantageous mode of operation is for such a module to be incorporated with the web server of a vendor in order to appropriately determine required level of trust for transactions initiated with the vendor's web server. In this way, transaction requests from users may be assigned a required trust level in accordance with the policies of the vendor and such information may be forwarded to the trust engine 110 along with the authentication request.

This required trust level correlates with the degree of certainty that the vendor wants to have that the individual authenticating is in fact who he identifies himself as. For example, if the transaction is one where the vendor wants a fair degree of certainty because goods are changing hands, the vendor may require a trust level of 85%. For situation where the vendor is merely authenticating the user to allow him to view members only content or exercise privileges on a chat room, the downside risk may be small enough that the vendor requires only a 60% trust level. However, to enter into a production contract with a value of tens of thousands of dollars, the vendor may require a trust level of 99% or more.

This required trust level represents a metric to which the user must authenticate himself in order to complete the transaction. If the required trust level is 85% for example, the user

must provide authentication to the trust engine 110 sufficient for the trust engine 110 to say with 85% confidence that the user is who they say they are. It is the balance between this required trust level and the authentication confidence level which produces either a positive authentication (to the satisfaction of the vendor) or a possibility of trust arbitrage.

As shown in FIG. 17, after the transaction engine 205 receives the required trust level, it compares in step 1720 the required trust level to the authentication confidence level which the authentication engine 215 calculated for the current authentication (as discussed with reference to FIG. 16). If the authentication confidence level is higher than the required trust level for the transaction in step 1730, then the process moves to step 1740 where a positive authentication for this transaction is produced by the transaction engine 205. A message to this effect will then be inserted into the authentication results and returned to the vendor by the transaction engine 205 as shown in step 1055 (see FIG. 10).

However, if the authentication confidence level does not fulfill the required trust level in step 1730, then a confidence gap exists for the current authentication, and trust arbitrage is conducted in step 1750. Trust arbitrage is described more completely with reference to FIG. 18 below. This process as described below takes place within the transaction engine 205 of the trust engine 110. Because no authentication or other cryptographic operations are needed to execute trust arbitrage (other than those required for the SSL communication between the transaction engine 205 and other components), the process may be performed outside the authentication engine 215. However, as will be discussed below, any reevaluation of authentication data or other cryptographic or authentication events will require the transaction engine 205 to resubmit the appropriate data to the authentication engine 215. Those of skill in the art will recognize that the trust arbitrage process could alternately be structured to take place partially or entirely within the authentication engine 215 itself.

As mentioned above, trust arbitrage is a process where the trust engine 110 mediates a negotiation between the vendor and user in an attempt to secure a positive authentication where appropriate. As shown in step 1805, the transaction engine 205 first determines whether or not the current situation is appropriate for trust arbitrage. This may be determined based upon the circumstances of the authentication, e.g. whether this authentication has already been through multiple cycles of arbitrage, as well as upon the preferences of either the vendor or user, as will be discussed further below.

In such circumstances where arbitrage is not possible, the process proceeds to step 1810 where the transaction engine 205 generates a negative authentication and then inserts it into the authentication results which are sent to the vendor in step 1055 (see FIG. 10). One limit which may be advantageously used to prevent authentications from pending indefinitely is to set a time-out period from the initial authentication request. In this way, any transaction which is not positively authenticated within the time limit is denied further arbitrage and negatively authenticated. Those of skill in the art will recognize that such a time limit may vary depending upon the circumstances of the transaction and the desires of the user and vendor. Limitations may also be placed upon the number of attempts that may be made at providing a successful authentication. Such limitations may be handled by an attempt limiter 535 as shown in FIG. 5.

If arbitrage is not prohibited in step 1805, the transaction engine 205 will then engage in negotiation with one or both of the transacting parties. The transaction engine 205 may send a message to the user requesting some form of additional

authentication in order to boost the authentication confidence level produced as shown in step 1820. In the simplest form, this may simply indicate that authentication was insufficient. A request to produce one or more additional authentication instances to improve the overall confidence level of the authentication may also be sent.

If the user provides some additional authentication instances in step 1825, then the transaction engine 205 adds these authentication instances to the authentication data for the transaction and forwards it to the authentication engine 215 as shown in step 1015 (see FIG. 10), and the authentication is reevaluated based upon both the pre-existing authentication instances for this transaction and the newly provided authentication instances.

An additional type of authentication may be a request from the trust engine 110 to make some form of person-to-person contact between the trust engine 110 operator (or a trusted associate) and the user, for example, by phone call. This phone call or other non-computer authentication can be used to provide personal contact with the individual and also to conduct some form of questionnaire based authentication. This also may give the opportunity to verify an originating telephone number and potentially a voice analysis of the user when he calls in. Even if no additional authentication data can be provided, the additional context associated with the user's phone number may improve the reliability of the authentication context. Any revised data or circumstances based upon this phone call are fed into the trust engine 110 for use in consideration of the authentication request.

Additionally, in step 1820 the trust engine 110 may provide an opportunity for the user to purchase insurance, effectively buying a more confident authentication. The operator of the trust engine 110 may, at times, only want to make such an option available if the confidence level of the authentication is above a certain threshold to begin with. In effect, this user side insurance is a way for the trust engine 110 to vouch for the user when the authentication meets the normal required trust level of the trust engine 110 for authentication, but does not meet the required trust level of the vendor for this transaction. In this way, the user may still successfully authenticate to a very high level as may be required by the vendor, even though he only has authentication instances which produce confidence sufficient for the trust engine 110.

This function of the trust engine 110 allows the trust engine 110 to vouch for someone who is authenticated to the satisfaction of the trust engine 110, but not of the vendor. This is analogous to the function performed by a notary in adding his signature to a document in order to indicate to someone reading the document at a later time that the person whose signature appears on the document is in fact the person who signed it. The signature of the notary testifies to the act of signing by the user. In the same way, the trust engine is providing an indication that the person transacting is who they say they are.

However, because the trust engine 110 is artificially boosting the level of confidence provided by the user, there is a greater risk to the trust engine 110 operator, since the user is not actually meeting the required trust level of the vendor. The cost of the insurance is designed to offset the risk of a false positive authentication to the trust engine 110 (who may be effectively notarizing the authentications of the user). The user pays the trust engine 110 operator to take the risk of authenticating to a higher level of confidence than has actually been provided.

Because such an insurance system allows someone to effectively buy a higher confidence rating from the trust engine 110, both vendors and users may wish to prevent the

use of user side insurance in certain transactions. Vendors may wish to limit positive authentications to circumstances where they know that actual authentication data supports the degree of confidence which they require and so may indicate to the trust engine **110** that user side insurance is not to be allowed. Similarly, to protect his online identity, a user may wish to prevent the use of user side insurance on his account, or may wish to limit its use to situations where the authentication confidence level without the insurance is higher than a certain limit. This may be used as a security measure to prevent someone from overhearing a password or stealing a smart card and using them to falsely authenticate to a low level of confidence, and then purchasing insurance to produce a very high level of (false) confidence. These factors may be evaluated in determining whether user side insurance is allowed.

If user purchases insurance in step **1840**, then the authentication confidence level is adjusted based upon the insurance purchased in step **1845**, and the authentication confidence level and required trust level are again compared in step **1730** (see FIG. **17**). The process continues from there, and may lead to either a positive authentication in step **1740** (see FIG. **17**), or back into the trust arbitrage process in step **1750** for either further arbitrage (if allowed) or a negative authentication in step **1810** if further arbitrage is prohibited.

In addition to sending a message to the user in step **1820**, the transaction engine **205** may also send a message to the vendor in step **1830** which indicates that a pending authentication is currently below the required trust level. The message may also offer various options on how to proceed to the vendor. One of these Options is to simply inform the vendor of what the current authentication confidence level is and ask if the vendor wishes to maintain their current unfulfilled required trust level. This may be beneficial because in some cases, the vendor may have independent means for authenticating the transaction or may have been using a default set of requirements which generally result in a higher required level being initially specified than is actually needed for the particular transaction at hand.

For instance, it may be standard practice that all incoming purchase order transactions with the vendor are expected to meet a 98% trust level. However, if an order was recently discussed by phone between the vendor and a long-standing customer, and immediately thereafter the transaction is authenticated, but only to a 93% confidence level, the vendor may wish to simply lower the acceptance threshold for this transaction, because the phone call effectively provides additional authentication to the vendor. In certain circumstances, the vendor may be willing to lower their required trust level, but not all the way to the level of the current authentication confidence. For instance, the vendor in the above example might consider that the phone call prior to the order might merit a 4% reduction in the degree of trust needed; however, this is still greater than the 93% confidence produced by the user.

If the vendor does adjust their required trust level in step **1835**, then the authentication confidence level produced by the authentication and the required trust level are compared in step **1730** (see FIG. **17**). If the confidence level now exceeds the required trust level, a positive authentication may be generated in the transaction engine **205** in step **1740** (see FIG. **17**). If not, further arbitrage may be attempted as discussed above if it is permitted.

In addition to requesting an adjustment to the required trust level, the transaction engine **205** may also offer vendor side insurance to the vendor requesting the authentication. This insurance serves a similar purpose to that described above for

the user side insurance. Here, however, rather than the cost corresponding to the risk being taken by the trust engine **110** in authenticating above the actual authentication confidence level produced, the cost of the insurance corresponds to the risk being taken by the vendor in accepting a lower trust level in the authentication.

Instead of just lowering their actual required trust level, the vendor has the option of purchasing insurance to protect itself from the additional risk associated with a lower level of trust in the authentication of the user. As described above, it may be advantageous for the vendor to only consider purchasing such insurance to cover the trust gap in conditions where the existing authentication is already above a certain threshold.

The availability of such vendor side insurance allows the vendor the option to either: lower his trust requirement directly at no additional cost to himself, bearing the risk of a false authentication himself (based on the lower trust level required); or, buying insurance for the trust gap between the authentication confidence level and his requirement, with the trust engine **110** operator bearing the risk of the lower confidence level which has been provided. By purchasing the insurance, the vendor effectively keeps his high trust level requirement; because the risk of a false authentication is shifted to the trust engine **110** operator.

If the vendor purchases insurance in step **1840**, the authentication confidence level and required trust levels are compared in step **1730** (see FIG. **17**), and the process continues as described above.

Note that it is also possible that both the user and the vendor respond to messages from the trust engine **110**. Those of skill in the art will recognize that there are multiple ways in which such situations can be handled. One advantageous mode of handling the possibility of multiple responses is simply to treat the responses in a first-come, first-served manner. For example, if the vendor responds with a lowered required trust level and immediately thereafter the user also purchases insurance to raise his authentication level, the authentication is first reevaluated based upon the lowered trust requirement from the vendor. If the authentication is now positive, the user's insurance purchase is ignored. In another advantageous mode of operation, the user might only be charged for the level of insurance required to meet the new, lowered trust requirement of the vendor (if a trust gap remained even with the lowered vendor trust requirement).

If no response from either party is received during the trust arbitrage process at step **1850** within the time limit set for the authentication, the arbitrage is reevaluated in step **1805**. This effectively begins the arbitrage process again. If the time limit was final or other circumstances prevent further arbitrage in step **1805**, a negative authentication is generated by the transaction engine **205** in step **1810** and returned to the vendor in step **1055** (see FIG. **10**). If not, new messages may be sent to the user and vendor, and the process may be repeated as desired.

Note that for certain types of transactions, for instance, digitally signing documents which are not part of a transaction, there may not necessarily be a vendor or other third party; therefore the transaction is primarily between the user and the trust engine **110**. In circumstances such as these, the trust engine **110** will have its own required trust level which must be satisfied in order to generate a positive authentication. However, in such circumstances, it will often not be desirable for the trust engine **110** to offer insurance to the user in order for him to raise the confidence of his own signature.

The process described above and shown in FIGS. **16-18** may be carried out using various communications modes as described above with reference to the trust engine **110**. For

instance, the messages may be web-based and sent using SSL connections between the trust engine 110 and applets downloaded in real time to browsers running on the user or vendor systems. In an alternate mode of operation, certain dedicated applications may be in use by the user and vendor which facilitate such arbitrage and insurance transactions. In another alternate mode of operation, secure email operations may be used to mediate the arbitrage described above, thereby allowing deferred evaluations and batch processing of authentications. Those of skill in the art will recognize that different communications modes may be used as are appropriate for the circumstances and authentication requirements of the vendor.

The following description with reference to FIG. 19 describes a sample transaction which integrates the various aspects of the present invention as described above. This example illustrates the overall process between a user and a vendor as mediated by the trust engine 110. Although the various steps and components as described in detail above may be used to carry out the following transaction, the process illustrated focuses on the interaction between the trust engine 110, user and vendor.

The transaction begins when the user, while viewing web pages online, fills out an order form on the web site of the vendor in step 1900. The user wishes to submit this order form to the vendor, signed with his digital signature. In order to do this, the user submits the order form with his request for a signature to the trust engine 110 in step 1905. The user will also provide authentication data which will be used as described above to authenticate his identity.

In step 1910 the authentication data is compared to the enrollment data by the trust engine 110 as discussed above, and if a positive authentication is produced, the hash of the order form, signed with the private key of the user, is forwarded to the vendor along with the order form itself.

The vendor receives the signed form in step 1915, and then the vendor will generate an invoice or other contract related to the purchase to be made in step 1920. This contract is sent back to the user with a request for a signature in step 1925. The vendor also sends an authentication request for this contract transaction to the trust engine 110 in step 1930 including a hash of the contract which will be signed by both parties. To allow the contract to be digitally signed by both parties, the vendor also includes authentication data for itself so that the vendor's signature upon the contract can later be verified if necessary.

As discussed above, the trust engine 110 then verifies the authentication data provided by the vendor to confirm the vendor's identity, and if the data produces a positive authentication in step 1935, continues with step 1955 when the data is received from the user. If the vendor's authentication data does not match the enrollment data of the vendor to the desired degree, a message is returned to the vendor requesting further authentication. Trust arbitrage may be performed here if necessary, as described above, in order for the vendor to successfully authenticate itself to the trust engine 110.

When the user receives the contract in step 1940, he reviews it, generates authentication data to sign it if it is acceptable in step 1945, and then sends a hash of the contract and his authentication data to the trust engine 110 in step 1950. The trust engine 110 verifies the authentication data in step 1955 and if the authentication is good, proceeds to process the contract as described below. As discussed above with reference to FIGS. 17 and 18, trust arbitrage may be performed as appropriate to close any trust gap which exists between the authentication confidence level and the required authentication level for the transaction.

The trust engine 110 signs the hash of the contract with the user's private key, and sends this signed hash to the vendor in step 1960, signing the complete message on its own behalf, i.e., including a hash of the complete message (including the user's signature) encrypted with the private key 510 of the trust engine 110. This message is received by the vendor in step 1965. The message represents a signed contract (hash of contract encrypted using user's private key) and a receipt from the trust engine 110 (the hash of the message including the signed contract, encrypted using the trust engine 110's private key).

The trust engine 110 similarly prepares a hash of the contract with the vendor's private key in step 1970, and forwards this to the user, signed by the trust engine 110. In this way, the user also receives a copy of the contract, signed by the vendor, as well as a receipt, signed by the trust engine 110, for delivery of the signed contract in step 1975.

In addition to the foregoing, an additional aspect of the invention provides a cryptographic Service Provider Module (SPM) which may be available to a client side application as a means to access functions provided by the trust engine 110 described above. One advantageous way to provide such a service is for the cryptographic SPM is to mediate communications between a third party Application Programming Interface (API) and a trust engine 110 which is accessible via a network or other remote connection. A sample cryptographic SPM is described below with reference to FIG. 20.

For example, on a typical system, a number of API's are available to programmers. Each API provides a set of function calls which may be made by an application 2000 running upon the system. Examples of API's which provide programming interfaces suitable for cryptographic functions, authentication functions, and other security function include the Cryptographic API (CAPI) 2010 provided by Microsoft with its Windows operating systems, and the Common Data Security Architecture (CDSA), sponsored by IBM, Intel and other members of the Open Group. CAPI will be used as an exemplary security API in the discussion that follows. However, the cryptographic SPM described could be used with CDSA or other security API's as are known in the art.

This API is used by a user system 105 or vendor system 120 when a call is made for a cryptographic function. Included among these functions may be requests associated with performing various cryptographic operations, such as encrypting a document with a particular key, signing a document, requesting a digital certificate, verifying a signature upon a signed document, and such other cryptographic functions as are described herein or known to those of skill in the art.

Such cryptographic functions are normally performed locally to the system upon which CAPI 2010 is located. This is because generally the functions called require the use of either resources of the local user system 105, such as a fingerprint reader, or software functions which are programmed using libraries which are executed on the local machine. Access to these local resources is normally provided by one or more Service Provider Modules (SPM's) 2015, 2020 as referred to above which provide resources with which the cryptographic functions are carried out. Such SPM's may include software libraries 2015 to perform encrypting or decrypting operations, or drivers and applications 2020 which are capable of accessing specialized hardware 2025, such as biometric scanning devices. In much the way that CAPI 2010 provides functions which may be used by an application 2000 of the system 105, the SPM's 2015, 2020 provide CAPI with access to the lower level functions and resources associated with the available services upon the system.

In accordance with the invention, it is possible to provide a cryptographic SPM 2030 which is capable of accessing the cryptographic functions provided by the trust engine 110 and making these functions available to an application 2000 through CAPI 2010. Unlike embodiments where CAPI 2010 is only able to access resources which are locally available through SPM's 2015, 2020, a cryptographic SPM 2030 as described herein would be able to submit requests for cryptographic operations to a remotely-located, network-accessible trust engine 110 in order to perform the operations desired.

For instance, if an application 2000 has a need for a cryptographic operation, such as signing a document, the application 2000 makes a function call to the appropriate CAPI 2010 function. CAPI 2010 in turn will execute this function, making use of the resources which are made available to it by the SPM's 2015, 2020 and the cryptographic SPM 2030. In the case of a digital signature function, the cryptographic SPM 2030 will generate an appropriate request which will be sent to the trust engine 110 across the communication link 125.

The operations which occur between the cryptographic SPM 2030 and the trust engine 110 are the same operations that would be possible between any other system and the trust engine 110. However, these functions are effectively made available to a user system 105 through CAPI 2010 such that they appear to be locally available upon the user system 105 itself. However, unlike ordinary SPM's 2015, 2020, the functions are being carried out on the remote trust engine 110 and the results relayed to the cryptographic SPM 2030 in response to appropriate requests across the communication link 125.

This cryptographic SPM 2030 makes a number of operations available to the user system 105 or a vendor system 120 which might not otherwise be available. These functions include without limitation: encryption and decryption of documents; issuance of digital certificates; digital signing of documents; verification of digital signatures; and such other operations as will be apparent to those of skill in the art.

In a separate embodiment, the present invention comprises a complete system for performing the data securing methods of the present invention on any data set. The computer system of this embodiment comprises a data splitting module that comprises the functionality shown in FIG. 8 and described herein. In one embodiment of the present invention, the data splitting module, sometimes referred to herein as a secure data parser, comprises a parser program or software suite which comprises data splitting, encryption and decryption, reconstitution or reassembly functionality. This embodiment may further comprise a data storage facility or multiple data storage facilities, as well. The data splitting module, or secure data parser, comprises a cross-platform software module suite which integrates within an electronic infrastructure, or as an add-on to any application which requires the ultimate security of its data elements. This parsing process operates on any type of data set, and on any and all file types, or in a database on any row, column or cell of data in that database.

The parsing process of the present invention may, in one embodiment, be designed in a modular tiered fashion, and any encryption process is suitable for use in the process of the present invention. The modular tiers of the parsing and splitting process of the present invention may include, but are not limited to, 1) cryptographic split, dispersed and securely stored in multiple locations; 2) encrypt, cryptographically split, dispersed and securely stored in multiple locations; 3) encrypt, cryptographically split, encrypt each share, then dispersed and securely stored in multiple locations; and 4) encrypt, cryptographically split, encrypt each share with a

different type of encryption than was used in the first step, then dispersed and securely stored in multiple locations.

The process comprises, in one embodiment, splitting of the data according to the contents of a generated random number, or key and performing the same cryptographic splitting of the key used in the encryption of splitting of the data to be secured into two or more portions, or shares, of parsed and split data, and in one embodiment, preferably into four or more portions of parsed and split data, encrypting all of the portions, then scattering and storing these portions back into the database, or relocating them to any named device, fixed or removable, depending on the requestor's need for privacy and security. Alternatively, in another embodiment, encryption may occur prior to the splitting of the data set by the splitting module or secure data parser. The original data processed as described in this embodiment is encrypted and obfuscated and is secured. The dispersion of the encrypted elements, if desired, can be virtually anywhere, including, but not limited to, a single server or data storage device, or among separate data storage facilities or devices. Encryption key management in one embodiment may be included within the software suite, or in another embodiment may be integrated into an existing infrastructure or any other desired location.

A cryptographic split (cryptosplit) partitions the data into N number of shares. The partitioning can be on any size unit of data, including an individual bit, bits, bytes, kilobytes, megabytes, or larger units, as well as any pattern or combination of data unit sizes whether predetermined or randomly generated. The units can also be of different sized, based on either a random or predetermined set of values. This means the data can be viewed as a sequence of these units. In this manner the size of the data units themselves may render the data more secure, for example by using one or more predetermined or randomly generated pattern, sequence or combination of data unit sizes. The units are then distributed (either randomly or by a predetermined set of values) into the N shares. This distribution could also involve a shuffling of the order of the units in the shares. It is readily apparent to those of ordinary skill in the art that the distribution of the data units into the shares may be performed according to a wide variety of possible selections, including but not limited to size-fixed, predetermined sizes, or one or more combination, pattern or sequence of data unit sizes that are predetermined or randomly generated.

In some embodiments of this cryptosplit split process, the data may be any suitable number of bytes in size, such as one, two, three, five, twenty, fifty, one hundred, more than one hundred, or N bytes in size. One particular example of this cryptographic split process, or cryptosplit, would be to consider the data to be 23 bytes in size, with the data unit size chosen to be one byte, and with the number of shares selected to be 4. Each byte would be distributed into one of the 4 shares. Assuming a random distribution, a key would be obtained to create a sequence of 23 random numbers (r1, r2, r3 through r23), each with a value between 1 and 4 corresponding to the four shares. Each of the units of data (in this example 23 individual bytes of data) is associated with one of the 23 random numbers corresponding to one of the four shares. The distribution of the bytes of data into the four shares would occur by placing the first byte of the data into share number r1, byte two into share r2, byte three into share r3, through the 23rd byte of data into share r23. It is readily apparent to those of ordinary skill in the art that a wide variety of other possible steps or combination or sequence of steps, including the size of the data units, may be used in the cryptosplit process of the present invention, and the above example is a non-limiting description of one process for

cryptosplitting data. To recreate the original data, the reverse operation would be performed.

In another embodiment of the cryptosplit process of the present invention, an option for the cryptosplitting process is to provide sufficient redundancy in the shares such that only a subset of the shares are needed to reassemble or restore the data to its original or useable form. As a non-limiting example, the cryptosplit may be done as a “3 of 4” cryptosplit such that only three of the four shares are necessary to reassemble or restore the data to its original or useable form. This is also referred to as a “M of N cryptosplit” wherein N is the total number of shares, and M is at least one less than N. It is readily apparent to those of ordinary skill in the art that there are many possibilities for creating this redundancy in the cryptosplitting process of the present invention.

In one embodiment of the cryptosplitting process of the present invention, each unit of data is stored in two shares, the primary share and the backup share. Using the “3 of 4” cryptosplitting process described above, any one share can be missing, and this is sufficient to reassemble or restore the original data with no missing data units since only three of the total four shares are required. As described herein, a random number is generated that corresponds to one of the shares. The random number is associated with a data unit, and stored in the corresponding share, based on a key. One key is used, in this embodiment, to generate the primary and backup share random number. As described herein for the cryptosplitting process of the present invention, a set of random numbers (also referred to as primary share numbers) from 0 to 3 are generated equal to the number of data units. Then another set of random numbers is generated (also referred to as backup share numbers) from 1 to 3 equal to the number of data units. Each unit of data is then associated with a primary share number and a backup share number. Alternatively, a set of random numbers may be generated that is fewer than the number of data units, and repeating the random number set, but this may reduce the security of the sensitive data. The primary share number is used to determine into which share the data unit is stored. The backup share number is combined with the primary share number to create a third share number between 0 and 3, and this number is used to determine into which share the data unit is stored. In this example, the equation to determine the third share number is:

$$\text{(primary share number+backup share number)MOD 4=third share number.}$$

In the embodiment described above where the primary share number is between 0 and 3, and the backup share number is between 1 and 3 ensures that the third share number is different from the primary share number. This results in the data unit being stored in two different shares. It is readily apparent to those of ordinary skill in the art that there are many ways of performing redundant cryptosplitting and non-redundant cryptosplitting in addition to the embodiments disclosed herein. For example, the data units in each share could be shuffled utilizing a different algorithm. This data unit shuffling may be performed as the original data is split into the data units, or after the data units are placed into the shares, or after the share is full, for example.

The various cryptosplitting processes and data shuffling processes described herein, and all other embodiments of the cryptosplitting and data shuffling methods of the present invention may be performed on data units of any size, including but not limited to, as small as an individual bit, bits, bytes, kilobytes, megabytes or larger.

An example of one embodiment of source code that would perform the cryptosplitting process described herein is:

```

DATA [1:24] - array of bytes with the data to be split
SHARES[0:3; 1:24] - 2-dimensional array with each row representing
one of the shares
5 RANDOM[1:24] - array random numbers in the range of 0..3
S1 = 1;
S2 = 1;
S3 = 1;
S4 = 1;
For J = 1 to 24 do
10 Begin
IF RANDOM[J] ==0 then
Begin
SHARES[1,S1] = DATA [J];
S1 = S1 + 1;
End
15 ELSE IF RANDOM[J] ==1 then
Begin
SHARES[2,S2] = DATA [J];
S2 = S2 + 1;
END
ELSE IF RANDOM[J] ==2 then
20 Begin
Shares[3,S3] = data [J];
S3 = S3 + 1;
End
Else begin
Shares[4,S4] = data [J];
S4 = S4 + 1;
25 End;
END;

```

An example of one embodiment of source code that would perform the cryptosplitting RAID process described herein is:

Generate two sets of numbers, PrimaryShare is 0 to 3, BackupShare is 1 to 3. Then put each data unit into share [primaryshare[1]] and share[(primaryshare[1]+backupshare [1]) mod 4, with the same process as in cryptosplitting described above. This method will be scalable to any size N, where only N-1 shares are necessary to restore the data.

The retrieval, recombining, reassembly or reconstituting of the encrypted data elements may utilize any number of authentication techniques, including, but not limited to, biometrics, such as fingerprint recognition, facial scan, hand scan, iris scan, retinal scan, ear scan, vascular pattern recognition or DNA analysis. The data splitting and/or parser modules of the present invention may be integrated into a wide variety of infrastructure products or applications as desired.

Traditional encryption technologies known in the art rely on one or more key used to encrypt the data and render it unusable without the key. The data, however, remains whole and intact and subject to attack. The secure data parser of the present invention, in one embodiment, addresses this problem by performing a cryptographic parsing and splitting of the encrypted file into two or more portions or shares, and in another embodiment, preferably four or more shares, adding another layer of encryption to each share of the data, then storing the shares in different physical and/or logical locations. When one or more data shares are physically removed from the system, either by using a removable device, such as a data storage device, or by placing the share under another party’s control, any possibility of compromise of secured data is effectively removed.

An example of one embodiment of the secure data parser of the present invention and an example of how it may be utilized is shown in FIG. 21 and described below. However, it is readily apparent to those of ordinary skill in the art that the secure data parser of the present invention may be utilized in a wide variety of ways in addition to the non-limiting example below. As a deployment option, and in one embodiment, the

secure data parser may be implemented with external session key management or secure internal storage of session keys. Upon implementation, a Parser Master Key will be generated which will be used for securing the application and for encryption purposes. It should be also noted that the incorporation of the Parser Master key in the resulting secured data allows for a flexibility of sharing of secured data by individuals within a workgroup, enterprise or extended audience.

As shown in FIG. 21, this embodiment of the present invention shows the steps of the process performed by the secure data parser on data to store the session master key with the parsed data:

1. Generating a session master key and encrypt the data using RS1 stream cipher.

2. Separating the resulting encrypted data into four shares or portions of parsed data according to the pattern of the session master key.

3. In this embodiment of the method, the session master key will be stored along with the secured data shares in a data depository. Separating the session master key according to the pattern of the Parser Master Key and append the key data to the encrypted parsed data.

4. The resulting four shares of data will contain encrypted portions of the original data and portions of the session master key. Generate a stream cipher key for each of the four data shares.

5. Encrypting each share, then store the encryption keys in different locations from the encrypted data portions or shares: Share 1 gets Key 4, Share 2 gets Key 1, Share 3 gets Key 2, Share 4 gets Key 3.

To restore the original data format, the steps are reversed.

It is readily apparent to those of ordinary skill in the art that certain steps of the methods described herein may be performed in different order, or repeated multiple times, as desired. It is also readily apparent to those skilled in the art that the portions of the data may be handled differently from one another. For example, multiple parsing steps may be performed on only one portion of the parsed data. Each portion of parsed data may be uniquely secured in any desirable way provided only that the data may be reassembled, reconstituted, reformed, decrypted or restored to its original or other usable form.

As shown in FIG. 22 and described herein, another embodiment of the present invention comprises the steps of the process performed by the secure data parser on data to store the session master key data in one or more separate key management table:

1. Generating a session master key and encrypt the data using RS1 stream cipher.

2. Separating the resulting encrypted data into four shares or portions of parsed data according to the pattern of the session master key.

3. In this embodiment of the method of the present invention, the session master key will be stored in a separate key management table in a data depository. Generating a unique transaction ID for this transaction. Storing the transaction ID and session master key in a separate key management table. Separating the transaction ID according to the pattern of the Parser Master Key and append the data to the encrypted parsed or separated data.

4. The resulting four shares of data will contain encrypted portions of the original data and portions of the transaction ID.

5. Generating a stream cipher key for each of the four data shares.

6. Encrypting each share, then store the encryption keys in different locations from the encrypted data portions or shares: Share 1 gets Key 4, Share 2 gets Key 1, Share 3 gets Key 2, Share 4 gets Key 3.

To restore the original data format, the steps are reversed.

It is readily apparent to those of ordinary skill in the art that certain steps of the method described herein may be performed in different order, or repeated multiple times, as desired. It is also readily apparent to those skilled in the art that the portions of the data may be handled differently from one another. For example, multiple separating or parsing steps may be performed on only one portion of the parsed data. Each portion of parsed data may be uniquely secured in any desirable way provided only that the data may be reassembled, reconstituted, reformed, decrypted or restored to its original or other usable form.

As shown in FIG. 23, this embodiment of the present invention shows the steps of the process performed by the secure data parser on data to store the session master key with the parsed data:

1. Accessing the parser master key associated with the authenticated user

2. Generating a unique Session Master key

3. Derive an Intermediary Key from an exclusive OR function of the Parser Master Key and Session Master key

4. Optional encryption of the data using an existing or new encryption algorithm keyed with the Intermediary Key.

5. Separating the resulting optionally encrypted data into four shares or portions of parsed data according to the pattern of the Intermediary key.

6. In this embodiment of the method, the session master key will be stored along with the secured data shares in a data depository. Separating the session master key according to the pattern of the Parser Master Key and append the key data to the optionally encrypted parsed data shares.

7. The resulting multiple shares of data will contain optionally encrypted portions of the original data and portions of the session master key.

8. Optionally generate an encryption key for each of the four data shares.

9. Optionally encrypting each share with an existing or new encryption algorithm, then store the encryption keys in different locations from the encrypted data portions or shares: for example, Share 1 gets Key 4, Share 2 gets Key 1, Share 3 gets Key 2, Share 4 gets Key 3.

To restore the original data format, the steps are reversed.

It is readily apparent to those of ordinary skill in the art that certain steps of the methods described herein may be performed in different order, or repeated multiple times, as desired. It is also readily apparent to those skilled in the art that the portions of the data may be handled differently from one another. For example, multiple parsing steps may be performed on only one portion of the parsed data. Each portion of parsed data may be uniquely secured in any desirable way provided only that the data may be reassembled, reconstituted, reformed, decrypted or restored to its original or other usable form.

As shown in FIG. 24 and described herein, another embodiment of the present invention comprises the steps of the process performed by the secure data parser on data to store the session master key data in one or more separate key management table:

1. Accessing the Parser Master Key associated with the authenticated user

2. Generating a unique Session Master Key

3. Derive an Intermediary Key from an exclusive OR function of the Parser Master Key and Session Master key

55

4. Optionally encrypt the data using an existing or new encryption algorithm keyed with the Intermediary Key.

5. Separating the resulting optionally encrypted data into four shares or portions of parsed data according to the pattern of the Intermediary Key.

6. In this embodiment of the method of the present invention, the session master key will be stored in a separate key management table in a data depository. Generating a unique transaction ID for this transaction. Storing the transaction ID and session master key in a separate key management table or passing the Session Master Key and transaction ID back to the calling program for external management. Separating the transaction ID according to the pattern of the Parser Master Key and append the data to the optionally encrypted parsed or separated data.

7. The resulting four shares of data will contain optionally encrypted portions of the original data and portions of the transaction ID.

8. Optionally generate an encryption key for each of the four data shares.

9. Optionally encrypting each share, then store the encryption keys in different locations from the encrypted data portions or shares. For example: Share 1 gets Key 4, Share 2 gets Key 1, Share 3 gets Key 2, Share 4 gets Key 3.

To restore the original data format, the steps are reversed.

It is readily apparent to those of ordinary skill in the art that certain steps of the method described herein may be performed in different order, or repeated multiple times, as desired. It is also readily apparent to those skilled in the art that the portions of the data may be handled differently from one another. For example, multiple separating or parsing steps may be performed on only one portion of the parsed data. Each portion of parsed data may be uniquely secured in any desirable way provided only that the data may be reassembled, reconstituted, reformed, decrypted or restored to its original or other usable form.

A wide variety of encryption methodologies are suitable for use in the methods of the present invention, as is readily apparent to those skilled in the art. The One Time Pad algorithm, is often considered one of the most secure encryption methods, and is suitable for use in the method of the present invention. Using the One Time Pad algorithm requires that a key be generated which is as long as the data to be secured. The use of this method may be less desirable in certain circumstances such as those resulting in the generation and management of very long keys because of the size of the data set to be secured. In the One-Time Pad (OTP) algorithm, the simple exclusive-or function, XOR, is used. For two binary streams x and y of the same length, x XOR y means the bitwise exclusive-or of x and y.

At the bit level is generated:

0 XOR 0=0

0 XOR 1=1

1 XOR 0=1

1 XOR 1=0

An example of this process is described herein for an n-byte secret, s, (or data set) to be split. The process will generate an n-byte random value, a, and then set:

$$b=a \text{ XOR } s.$$

Note that one can derive "s" via the equation:

$$s=a \text{ XOR } b.$$

The values a and b are referred to as shares or portions and are placed in separate depositories. Once the secret s is split into two or more shares, it is discarded in a secure manner.

56

The secure data parser of the present invention may utilize this function, performing multiple XOR functions incorporating multiple distinct secret key values: K1, K2, K3, Kn, K5. At the beginning of the operation, the data to be secured is passed through the first encryption operation, secure data=data XOR secret key 5:

$$S=D \text{ XOR } K5$$

In order to securely store the resulting encrypted data in, for example, four shares, S1, S2, S3, Sn, the data is parsed and split into "n" segments, or shares, according to the value of K5. This operation results in "n" pseudorandom shares of the original encrypted data. Subsequent XOR functions may then be performed on each share with the remaining secret key values, for example: Secure data segment 1=encrypted data share 1 XOR secret key 1:

$$SD1=S1 \text{ XOR } K1$$

$$SD2=S2 \text{ XOR } K2$$

$$SD3=S3 \text{ XOR } K3$$

$$SDn=Sn \text{ XOR } Kn.$$

In one embodiment, it may not be desired to have any one depository contain enough information to decrypt the information held there, so the key required to decrypt the share is stored in a different data depository:

Depository 1: SD1, Kn

Depository 2: SD2, K1

Depository 3: SD3, K2

Depository n: SDn, K3.

Additionally, appended to each share may be the information required to retrieve the original session encryption key, K5. Therefore, in the key management example described herein, the original session master key is referenced by a transaction ID split into "n" shares according to the contents of the installation dependant Parser Master Key (TID1, TID2, TID3, TIDn):

Depository 1: SD1, Kn, TID1

Depository 2: SD2, K1, TID2

Depository 3: SD3, K2, TID3

Depository n: SDn, K3, TIDn.

In the incorporated session key example described herein, the session master key is split into "n" shares according to the contents of the installation dependant Parser Master Key (SK1, SK2, SK3, SKn):

Depository 1: SD1, Kn, SK1

Depository 2: SD2, K1, SK2

Depository 3: SD3, K2, SK3

Depository n: SDn, K3, SKn.

Unless all four shares are retrieved, the data cannot be reassembled according to this example. Even if all four shares are captured, there is no possibility of reassembling or restoring the original information without access to the session master key and the Parser Master Key.

This example has described an embodiment of the method of the present invention, and also describes, in another embodiment, the algorithm used to place shares into depositories so that shares from all depositories can be combined to form the secret authentication material. The computations needed are very simple and fast. However, with the One Time Pad (OTP) algorithm there may be circumstances that cause it to be less desirable, such as a large data set to be secured, because the key size is the same size as the data to be stored. Therefore, there would be a need to store and transmit about twice the amount of the original data which may be less desirable under certain circumstances.

Stream Cipher RS1

The stream cipher RS1 splitting technique is very similar to the OTP splitting technique described herein. Instead of an

n-byte random value, an $n' = \min(n, 16)$ —byte random value is generated and used to key the RS1 Stream Cipher algorithm. The advantage of the RS1 Stream Cipher algorithm is that a pseudorandom key is generated from a much smaller seed number. The speed of execution of the RS1 Stream Cipher encryption is also rated at approximately 10 times the speed of the well known in the art Triple DES encryption without compromising security. The RS1 Stream Cipher algorithm is well known in the art, and may be used to generate the keys used in the XOR function. The RS1 Stream Cipher algorithm is interoperable with other commercially available stream cipher algorithms, such as the RC4™ stream cipher algorithm of RSA Security, Inc and is suitable for use in the methods of the present invention.

Using the key notation above, K1 thru K5 are now an n' byte random values and we set:

$SD1 = S1 \text{ XOR } E(K1)$

$SD2 = S2 \text{ XOR } E(K2)$

$SD3 = S3 \text{ XOR } E(K3)$

$SDn = Sn \text{ XOR } E(Kn)$

where $E(K1)$ thru $E(Kn)$ are the first n' bytes of output from the RS1 Stream Cipher algorithm keyed by K1 thru Kn. The shares are now placed into data depositories as described herein.

In this stream cipher RS1 algorithm, the required computations needed are nearly as simple and fast as the OTP algorithm. The benefit in this example using the RS1 Stream Cipher is that the system needs to store and transmit on average only about 16 bytes more than the size of the original data to be secured per share. When the size of the original data is more than 16 bytes, this RS1 algorithm is more efficient than the OTP algorithm because it is simply shorter. It is readily apparent to those of ordinary skill in the art that a wide variety of encryption methods or algorithms are suitable for use in the present invention, including, but not limited to RS1, OTP, RC4™, Triple DES and AES.

There are major advantages provided by the data security methods and computer systems of the present invention over traditional encryption methods. One advantage is the security gained from moving shares of the data to different locations on one or more data depositories or storage devices, that may be in different logical, physical or geographical locations. When the shares of data are split physically and under the control of different personnel, for example, the possibility of compromising the data is greatly reduced.

Another advantage provided by the methods and system of the present invention is the combination of the steps of the method of the present invention for securing data to provide a comprehensive process of maintaining security of sensitive data. The data is encrypted with a secure key and split into one or more shares, and in one embodiment, four shares, according to the secure key. The secure key is stored safely with a reference pointer which is secured into four shares according to a secure key. The data shares are then encrypted individually and the keys are stored safely with different encrypted shares. When combined, the entire process for securing data according to the methods disclosed herein becomes a comprehensive package for data security.

The data secured according to the methods of the present invention is readily retrievable and restored, reconstituted, reassembled, decrypted, or otherwise returned into its original or other suitable form for use. In order to restore the original data, the following items may be utilized:

1. All shares or portions of the data set.
2. Knowledge of and ability to reproduce the process flow of the method used to secure the data.
3. Access to the session master key.

4. Access to the Parser Master Key.

Therefore, it may be desirable to plan a secure installation wherein at least one of the above elements may be physically separated from the remaining components of the system (under the control of a different system administrator for example).

Protection against a rogue application invoking the data securing methods application may be enforced by use of the Parser Master Key. A mutual authentication handshake between the secure data parser and the application may be required in this embodiment of the present invention prior to any action taken.

The security of the system dictates that there be no “back-door” method for recreation of the original data. For installations where data recovery issues may arise, the secure data parser can be enhanced to provide a minor of the four shares and session master key depository. Hardware options such as RAID (redundant array of inexpensive disks, used to spread information over several disks) and software options such as replication can assist as well in the data recovery planning.

Key Management

In one embodiment of the present invention, the data securing method uses three sets of keys for an encryption operation. Each set of keys may have individual key storage, retrieval, security and recovery options, based on the installation. The keys that may be used, include, but are not limited to:

The Parser Master Key

This key is an individual key associated with the installation of the secure data parser. It is installed on the server on which the secure data parser has been deployed. There are a variety of options suitable for securing this key including, but not limited to, a smart card, separate hardware key store, standard key stores, custom key stores or within a secured database table, for example.

The Session Master Key

A Session Master Key may be generated each time data is secured. The Session Master Key is used to encrypt the data prior to the parsing and splitting operations. It may also be incorporated (if the Session Master Key is not integrated into the parsed data) as a means of parsing the encrypted data. The Session Master Key may be secured in a variety of manners, including, but not limited to, a standard key store, custom key store, separate database table, or secured within the encrypted shares, for example.

The Share Encryption Keys

For each share or portions of a data set that is created, an individual Share Encryption Key may be generated to further encrypt the shares. The Share Encryption Keys may be stored in different shares than the share that was encrypted.

It is readily apparent to those of ordinary skill in the art that the data securing methods and computer system of the present invention are widely applicable to any type of data in any setting or environment. In addition to commercial applications conducted over the Internet or between customers and vendors, the data securing methods and computer systems of the present invention are highly applicable to non-commercial or private settings or environments. Any data set that is desired to be kept secure from any unauthorized user may be secured using the methods and systems described herein. For example, access to a particular database within a company or organization may be advantageously restricted to only selected users by employing the methods and systems of the present invention for securing data. Another example is the generation, modification or access to documents wherein it is desired to restrict access or prevent unauthorized or accidental access or disclosure outside a group of selected individu-

als, computers or workstations. These and other examples of the ways in which the methods and systems of data securing of the present invention are applicable to any non-commercial or commercial environment or setting for any setting, including, but not limited to any organization, government agency or corporation.

In another embodiment of the present invention, the data securing method uses three sets of keys for an encryption operation. Each set of keys may have individual key storage, retrieval, security and recovery options, based on the installation. The keys that may be used, include, but are not limited to:

1. The Parser Master Key

This key is an individual key associated with the installation of the secure data parser. It is installed on the server on which the secure data parser has been deployed. There are a variety of options suitable for securing this key including, but not limited to, a smart card, separate hardware key store, standard key stores, custom key stores or within a secured database table, for example.

2. The Session Master Key

A Session Master Key may be generated each time data is secured. The Session Master Key is used in conjunction with the Parser Master key to derive the Intermediary Key. The Session Master Key may be secured in a variety of manners, including, but not limited to, a standard key store, custom key store, separate database table, or secured within the encrypted shares, for example.

3. The Intermediary Key

An Intermediary Key may be generated each time data is secured. The Intermediary Key is used to encrypt the data prior to the parsing and splitting operation. It may also be incorporated as a means of parsing the encrypted data.

4. The Share Encryption Keys

For each share or portions of a data set that is created, an individual Share Encryption Key may be generated to further encrypt the shares. The Share Encryption Keys may be stored in different shares than the share that was encrypted.

It is readily apparent to those of ordinary skill in the art that the data securing methods and computer system of the present invention are widely applicable to any type of data in any setting or environment. In addition to commercial applications conducted over the Internet or between customers and vendors, the data securing methods and computer systems of the present invention are highly applicable to non-commercial or private settings or environments. Any data set that is desired to be kept secure from any unauthorized user may be secured using the methods and systems described herein. For example, access to a particular database within a company or organization may be advantageously restricted to only selected users by employing the methods and systems of the present invention for securing data. Another example is the generation, modification or access to documents wherein it is desired to restrict access or prevent unauthorized or accidental access or disclosure outside a group of selected individuals, computers or workstations. These and other examples of the ways in which the methods and systems of data securing of the present invention are applicable to any non-commercial or commercial environment or setting for any setting, including, but not limited to any organization, government agency or corporation.

Workgroup, Project, Individual PC/Laptop or Cross Platform Data Security

The data securing methods and computer systems of the present invention are also useful in securing data by workgroup, project, individual PC/Laptop and any other platform that is in use in, for example, businesses, offices, government

agencies, or any setting in which sensitive data is created, handled or stored. The present invention provides methods and computer systems to secure data that is known to be sought after by organizations, such as the U.S. Government, for implementation across the entire government organization or between governments at a state or federal level.

The data securing methods and computer systems of the present invention provide the ability to not only parse and split flat files but also data fields, sets and or table of any type. Additionally, all forms of data are capable of being secured under this process, including, but not limited to, text, video, images, biometrics and voice data. Scalability, speed and data throughput of the methods of securing data of the present invention are only limited to the hardware the user has at their disposal.

In one embodiment of the present invention, the data securing methods are utilized as described below in a workgroup environment. In one embodiment, as shown in FIG. 23 and described below, the Workgroup Scale data securing method of the present invention uses the private key management functionality of the TrustEngine to store the user/group relationships and the associated private keys (Parser Group Master Keys) necessary for a group of users to share secure data. The method of the present invention has the capability to secure data for an enterprise, workgroup, or individual user, depending on how the Parser Master Key was deployed.

In one embodiment, additional key management and user/group management programs may be provided, enabling wide scale workgroup implementation with a single point of administration and key management. Key generation, management and revocation are handled by the single maintenance program, which all become especially important as the number of users increase. In another embodiment, key management may also be set up across one or several different system administrators, which may not allow any one person or group to control data as needed. This allows for the management of secured data to be obtained by roles, responsibilities, membership, rights, etc., as defined by an organization, and the access to secured data can be limited to just those who are permitted or required to have access only to the portion they are working on, while others, such as managers or executives, may have access to all of the secured data. This embodiment allows for the sharing of secured data among different groups within a company or organization while at the same time only allowing certain selected individuals, such as those with the authorized and predetermined roles and responsibilities, to observe the data as a whole. In addition, this embodiment of the methods and systems of the present invention also allows for the sharing of data among, for example, separate companies, or separate departments or divisions of companies, or any separate organization departments, groups, agencies, or offices, or the like, of any government or organization or any kind, where some sharing is required, but not any one party may be permitted to have access to all the data. Particularly apparent examples of the need and utility for such a method and system of the present invention are to allow sharing, but maintain security, in between government areas, agencies and offices, and between different divisions, departments or offices of a large company, or any other organization, for example.

An example of the applicability of the methods of the present invention on a smaller scale is as follows. A Parser Master key is used as a serialization or branding of the secure data parser to an organization. As the scale of use of the Parser Master key is reduced from the whole enterprise to a smaller workgroup, the data securing methods described herein are used to share files within groups of users.

61

In the example shown in FIG. 25 and described below, there are six users defined along with their title or role within the organization. The side bar represents five possible groups that the users can belong to according to their role. The arrow represents membership by the user in one or more of the groups.

When configuring the secure data parser for use in this example, the system administrator accesses the user and group information from the operating system by a maintenance program. This maintenance program generates and assigns Parser Group Master Keys to users based on their membership in groups.

In this example, there are three members in the Senior Staff group. For this group, the actions would be:

1. Access Parser Group Master Key for the Senior Staff group (generate a key if not available);
2. Generate a digital certificate associating CEO with the Senior Staff group;
3. Generate a digital certificate associating CFO with the Senior Staff group;
4. Generate a digital certificate associating Vice President, Marketing with the Senior Staff group.

The same set of actions would be done for each group, and each member within each group. When the maintenance program is complete, the Parser Group Master Key becomes a shared credential for each member of the group. Revocation of the assigned digital certificate may be done automatically when a user is removed from a group through the maintenance program without affecting the remaining members of the group.

Once the shared credentials have been defined, the parsing and splitting process remains the same. When a file, document or data element is to be secured, the user is prompted for the target group to be used when securing the data. The resulting secured data is only accessible by other members of the target group. This functionality of the methods and systems of the present invention may be used with any other computer system or software platform, any may be, for example, integrated into existing application programs or used standalone for file security.

It is readily apparent to those of ordinary skill in the art that any one or combination of encryption algorithms are suitable for use in the methods and systems of the present invention. For example, the encryption steps may, in one embodiment, be repeated to produce a multi-layered encryption scheme. In addition, a different encryption algorithm, or combination of encryption algorithms, may be used in repeat encryption steps such that different encryption algorithms are applied to the different layers of the multi-layered encryption scheme. As such, the encryption scheme itself may become a component of the methods of the present invention for securing sensitive data from unauthorized use or access.

The secure data parser may include as an internal component, as an external component, or as both an error-checking component. For example, in one suitable approach, as portions of data are created using the secure data parser in accordance with the present invention, to assure the integrity of the data within a portion, a hash value is taken at preset intervals within the portion and is appended to the end of the interval. The hash value is a predictable and reproducible numeric representation of the data. If any bit within the data changes, the hash value would be different. A scanning module (either as a stand-alone component external to the secure data parser or as an internal component) may then scan the portions of data generated by the secure data parser. Each portion of data (or alternatively, less than all portions of data according to some interval or by a random or pseudo-random sampling) is

62

compared to the appended hash value or values and an action may be taken. This action may include a report of values that match and do not match, an alert for values that do not match, or invoking of some external or internal program to trigger a recovery of the data. For example, recovery of the data could be performed by invoking a recovery module based on the concept that fewer than all portions may be needed to generate original data in accordance with the present invention.

Any other suitable integrity checking may be implemented using any suitable integrity information appended anywhere in all or a subset of data portions. Integrity information may include any suitable information that can be used to determine the integrity of data portions. Examples of integrity information may include hash values computed based on any suitable parameter (e.g., based on respective data portions), digital signature information, message authentication code (MAC) information, any other suitable information, or any combination thereof.

The secure data parser of the present invention may be used in any suitable application. Namely, the secure data parser described herein has a variety of applications in different areas of computing and technology. Several such areas are discussed below. It will be understood that these are merely illustrative in nature and that any other suitable applications may make use of the secure data parser. It will further be understood that the examples described are merely illustrative embodiments that may be modified in any suitable way in order to satisfy any suitable desires. For example, parsing and splitting may be based on any suitable units, such as by bits, by bytes, by kilobytes, by megabytes, by any combination thereof, or by any other suitable unit.

The secure data parser of the present invention may be used to implement secure physical tokens, whereby data stored in a physical token may be required in order to access additional data stored in another storage area. In one suitable approach, a physical token, such as a compact USB flash drive, a floppy disk, an optical disk, a smart card, or any other suitable physical token, may be used to store one of at least two portions of parsed data in accordance with the present invention. In order to access the original data, the USB flash drive would need to be accessed. Thus, a personal computer holding one portion of parsed data would need to have the USB flash drive, having the other portion of parsed data, attached before the original data can be accessed. FIG. 26 illustrates this application. Storage area 2500 includes a portion of parsed data 2502. Physical token 2504, having a portion of parsed data 2506 would need to be coupled to storage area 2500 using any suitable communications interface 2508 (e.g., USB, serial, parallel, Bluetooth, IR, IEEE 1394, Ethernet, or any other suitable communications interface) in order to access the original data. This is useful in a situation where, for example, sensitive data on a computer is left alone and subject to unauthorized access attempts. By removing the physical token (e.g., the USB flash drive), the sensitive data is inaccessible. It will be understood that any other suitable approach for using physical tokens may be used.

The secure data parser of the present invention may be used to implement a secure authentication system whereby user enrollment data (e.g., passwords, private encryption keys, fingerprint templates, biometric data or any other suitable user enrollment data) is parsed and split using the secure data parser. The user enrollment data may be parsed and split whereby one or more portions are stored on a smart card, a government Common Access Card, any suitable physical storage device (e.g., magnetic or optical disk, USB key drive, etc.), or any other suitable device. One or more other portions of the parsed user enrollment data may be stored in the system

63

performing the authentication. This provides an added level of security to the authentication process (e.g., in addition to the biometric authentication information obtained from the biometric source, the user enrollment data must also be obtained via the appropriate parsed and split data portion).

The secure data parser of the present invention may be integrated into any suitable existing system in order to provide the use of its functionality in each system's respective environment. FIG. 27 shows a block diagram of an illustrative system 2600, which may include software, hardware, or both for implementing any suitable application. System 2600 may be an existing system in which secure data parser 2602 may be retrofitted as an integrated component. Alternatively, secure data parser 2602 may be integrated into any suitable system 2600 from, for example, its earliest design stage. Secure data parser 2600 may be integrated at any suitable level of system 2600. For example, secure data parser 2602 may be integrated into system 2600 at a sufficiently back-end level such that the presence of secure data parser 2602 may be substantially transparent to an end user of system 2600. Secure data parser 2602 may be used for parsing and splitting data among one or more storage devices 2604 in accordance with the present invention. Some illustrative examples of systems having the secure data parser integrated therein are discussed below.

The secure data parser of the present invention may be integrated into an operating system kernel (e.g., Linux, Unix, or any other suitable commercial or proprietary operating system). This integration may be used to protect data at the device level whereby, for example, data that would ordinarily be stored in one or more devices is separated into a certain number of portions by the secure data parser integrated into the operating system and stored among the one or more devices. When original data is attempted to be accessed, the appropriate software, also integrated into the operating system, may recombine the parsed data portions into the original data in a way that may be transparent to the end user.

The secure data parser of the present invention may be integrated into a volume manager or any other suitable component of a storage system to protect local and networked data storage across any or all supported platforms. For example, with the secure data parser integrated, a storage system may make use of the redundancy offered by the secure data parser (i.e., which is used to implement the feature of needing fewer than all separated portions of data in order to reconstruct the original data) to protect against data loss. The secure data parser also allows all data written to storage devices, whether using redundancy or not, to be in the form of multiple portions that are generated according to the parsing of the present invention. When original data is attempted to be accessed, the appropriate software, also integrated into the volume manager or other suitable component of the storage system, may recombine the parsed data portions into the original data in a way that may be transparent to the end user.

In one suitable approach, the secure data parser of the present invention may be integrated into a RAID controller (as either hardware or software). This allows for the secure storage of data to multiple drives while maintaining fault tolerance in case of drive failure.

The secure data parser of the present invention may be integrated into a database in order to, for example, protect sensitive table information. For example, in one suitable approach, data associated with particular cells of a database table (e.g., individual cells, one or more particular columns, one or more particular rows, any combination thereof, or an entire database table) may be parsed and separated according to the present invention (e.g., where the different portions are

64

stored on one or more storage devices at one or more locations or on a single storage device). Access to recombine the portions in order to view the original data may be granted by traditional authentication methods (e.g., username and password query).

The secure data parser of the present invention may be integrated in any suitable system that involves data in motion (i.e., transfer of data from one location to another). Such systems include, for example, email, streaming data broadcasts, and wireless (e.g., WiFi) communications. With respect to email, in one suitable approach, the secure data parser may be used to parse outgoing messages (i.e., containing text, binary data, or both (e.g., files attached to an email message)) and sending the different portions of the parsed data along different paths thus creating multiple streams of data. If any one of these streams of data is compromised, the original message remains secure because the system may require that more than one of the portions be combined, in accordance with the present invention, in order to generate the original data. In another suitable approach, the different portions of data may be communicated along one path sequentially so that if one portion is obtained, it may not be sufficient to generate the original data. The different portions arrive at the intended recipient's location and may be combined to generate the original data in accordance with the present invention.

FIGS. 28 and 29 are illustrative block diagrams of such email systems. FIG. 28 shows a sender system 2700, which may include any suitable hardware, such as a computer terminal, personal computer, handheld device (e.g., PDA, Blackberry), cellular telephone, computer network, any other suitable hardware, or any combination thereof. Sender system 2700 is used to generate and/or store a message 2704, which may be, for example, an email message, a binary data file (e.g., graphics, voice, video, etc.), or both. Message 2704 is parsed and split by secure data parser 2702 in accordance with the present invention. The resultant data portions may be communicated across one or more separate communications paths 2706 over network 2708 (e.g., the Internet, an intranet, a LAN, WiFi, Bluetooth, any other suitable hard-wired or wireless communications means, or any combination thereof) to recipient system 2710. The data portions may be communicated parallel in time or alternatively, according to any suitable time delay between the communication of the different data portions. Recipient system 2710 may be any suitable hardware as described above with respect to sender system 2700. The separate data portions carried along communications paths 2706 are recombined at recipient system 2710 to generate the original message or data in accordance with the present invention.

FIG. 29 shows a sender system 2800, which may include any suitable hardware, such as a computer terminal, personal computer, handheld device (e.g., PDA), cellular telephone, computer network, any other suitable hardware, or any combination thereof. Sender system 2800 is used to generate and/or store a message 2804, which may be, for example, an email message, a binary data file (e.g., graphics, voice, video, etc.), or both. Message 2804 is parsed and split by secure data parser 2802 in accordance with the present invention. The resultant data portions may be communicated across a single communications paths 2806 over network 2808 (e.g., the Internet, an intranet, a LAN, WiFi, Bluetooth, any other suitable communications means, or any combination thereof) to recipient system 2810. The data portions may be communicated serially across communications path 2806 with respect to one another. Recipient system 2810 may be any suitable hardware as described above with respect to sender system 2800. The separate data portions carried along communica-

65

tions path 2806 are recombined at recipient system 2810 to generate the original message or data in accordance with the present invention.

It will be understood that the arrangement of FIGS. 28 and 29 are merely illustrative. Any other suitable arrangement may be used. For example, in another suitable approach, the features of the systems of FIGS. 28 and 29 may be combined whereby the multi-path approach of FIG. 28 is used and in which one or more of communications paths 2706 are used to carry more than one portion of data as communications path 2806 does in the context of FIG. 29.

The secure data parser may be integrated at any suitable level of a data-in motion system. For example, in the context of an email system, the secure data parser may be integrated at the user-interface level (e.g., into Microsoft® Outlook), in which case the user may have control over the use of the secure data parser features when using email. Alternatively, the secure data parser may be implemented in a back-end component such as at the exchange server, in which case messages may be automatically parsed, split, and communicated along different paths in accordance with the present invention without any user intervention.

Similarly, in the case of streaming broadcasts of data (e.g., audio, video), the outgoing data may be parsed and separated into multiple streams each containing a portion of the parsed data. The multiple streams may be transmitted along one or more paths and recombined at the recipient's location in accordance with the present invention. One of the benefits of this approach is that it avoids the relatively large overhead associated with traditional encryption of data followed by transmission of the encrypted data over a single communications channel. The secure data parser of the present invention allows data in motion to be sent in multiple parallel streams, increasing speed and efficiency.

It will be understood that the secure data parser may be integrated for protection of and fault tolerance of any type of data in motion through any transport medium, including, for example, wired, wireless, or physical. For example, voice over Internet protocol (VoIP) applications may make use of the secure data parser of the present invention. Wireless or wired data transport from or to any suitable personal digital assistant (PDA) devices such as Blackberries and Smart-Phones may be secured using the secure data parser of the present invention. Communications using wireless 802.11 protocols for peer to peer and hub based wireless networks, satellite communications, point to point wireless communications, Internet client/server communications, or any other suitable communications may involve the data in motion capabilities of the secure data parser in accordance with the present invention. Data communication between computer peripheral device (e.g., printer, scanner, monitor, keyboard, network router, biometric authentication device (e.g., fingerprint scanner), or any other suitable peripheral device) between a computer and a computer peripheral device, between a computer peripheral device and any other suitable device, or any combination thereof may make use of the data in motion features of the present invention.

The data in motion features of the present invention may also apply to physical transportation of secure shares using for example, separate routes, vehicles, methods, any other suitable physical transportation, or any combination thereof. For example, physical transportation of data may take place on digital/magnetic tapes, floppy disks, optical disks, physical tokens, USB drives, removable hard drives, consumer electronic devices with flash memory (e.g., Apple IPODs or other MP3 players), flash memory, any other suitable medium used for transporting data, or any combination thereof.

66

The secure data parser of the present invention may provide security with the ability for disaster recovery. According to the present invention, fewer than all portions of the separated data generated by the secure data parser may be necessary in order to retrieve the original data. That is, out of m portions stored, n may be the minimum number of these m portions necessary to retrieve the original data, where $n \leq m$. For example, if each of four portions is stored in a different physical location relative to the other three portions, then, if $n=2$ in this example, two of the locations may be compromised whereby data is destroyed or inaccessible, and the original data may still be retrieved from the portions in the other two locations. Any suitable value for n or m may be used.

In addition, the n of m feature of the present invention may be used to create a "two man rule" whereby in order to avoid entrusting a single individual or any other entity with full access to what may be sensitive data, two or more distinct entities, each with a portion of the separated data parsed by the secure data parser of the present invention may need to agree to put their portions together in order to retrieve the original data.

The secure data parser of the present invention may be used to provide a group of entities with a group-wide key that allows the group members to access particular information authorized to be accessed by that particular group. The group key may be one of the data portions generated by the secure data parser in accordance with the present invention that may be required to be combined with another portion centrally stored, for example in order to retrieve the information sought. This feature allows for, for example, secure collaboration among a group. It may be applied in for example, dedicated networks, virtual private networks, intranets, or any other suitable network.

Specific applications of this use of the secure data parser include, for example, coalition information sharing in which, for example, multi-national friendly government forces are given the capability to communicate operational and otherwise sensitive data on a security level authorized to each respective country over a single network or a dual network (i.e., as compared to the many networks involving relatively substantial manual processes currently used). This capability is also applicable for companies or other organizations in which information needed to be known by one or more specific individuals (within the organization or without) may be communicated over a single network without the need to worry about unauthorized individuals viewing the information.

Another specific application includes a multi-level security hierarchy for government systems. That is, the secure data parser of the present invention may provide for the ability to operate a government system at different levels of classified information (e.g., unclassified, classified, secret, top secret) using a single network. If desired, more networks may be used (e.g., a separate network for top secret), but the present invention allows for substantially fewer than current arrangement in which a separate network is used for each level of classification.

It will be understood that any combination of the above described applications of the secure data parser of the present invention may be used. For example, the group key application can be used together with the data in motion security application (i.e., whereby data that is communicated over a network can only be accessed by a member of the respective group and where, while the data is in motion, it is split among multiple paths (or sent in sequential portions) in accordance with the present invention).

The secure data parser of the present invention may be integrated into any middleware application to enable applications to securely store data to different database products or to different devices without modification to either the applications or the database. Middleware is a general term for any product that allows two separate and already existing programs to communicate. For example, in one suitable approach, middleware having the secure data parser integrated, may be used to allow programs written for a particular database to communicate with other databases without custom coding.

The secure data parser of the present invention may be implemented having any combination of any suitable capabilities, such as those discussed herein. In some embodiments of the present invention, for example, the secure data parser may be implemented having only certain capabilities whereas other capabilities may be obtained through the use of external software, hardware, or both interfaced either directly or indirectly with the secure data parser.

FIG. 30, for example, shows an illustrative implementation of the secure data parser as secure data parser 3000. Secure data parser 3000 may be implemented with very few built-in capabilities. As illustrated, secure data parser 3000 may include built-in capabilities for parsing and splitting data into portions (also referred to herein as shares) of data using module 3002 in accordance with the present invention. Secure data parser 3000 may also include built in capabilities for performing redundancy in order to be able to implement, for example, the m of n feature described above (i.e., recreating the original data using fewer than all shares of parsed and split data) using module 3004. Secure data parser 3000 may also include share distribution capabilities using module 3006 for placing the shares of data into buffers from which they are sent for communication to a remote location, for storage, etc. in accordance with the present invention. It will be understood that any other suitable capabilities may be built into secure data parser 3000.

Assembled data buffer 3008 may be any suitable memory used to store the original data (although not necessarily in its original form) that will be parsed and split by secure data parser 3000. In a splitting operation, assembled data buffer 3008 provides input to secure data parser 3000. In a restore operation, assembled data buffer 3008 may be used to store the output of secure data parser 3000.

Split shares buffers 3010 may be one or more memory modules that may be used to store the multiple shares of data that resulted from the parsing and splitting of original data. In a splitting operation, split shares buffers 3010 hold the output of the secure data parser. In a restore operation, split shares buffers hold the input to secure data parser 3000.

It will be understood that any other suitable arrangement of capabilities may be built-in for secure data parser 3000. Any additional features may be built-in and any of the features illustrated may be removed, made more robust, made less robust, or may otherwise be modified in any suitable way. Buffers 3008 and 3010 are likewise merely illustrative and may be modified, removed, or added to in any suitable way.

Any suitable modules implemented in software, hardware or both may be called by or may call to secure data parser 3000. If desired, even capabilities that are built into secure data parser 3000 may be replaced by one or more external modules. As illustrated, some external modules include random number generator 3012, cipher feedback key generator 3014, hash algorithm 3016, any one or more types of encryption 3018, and key management 3020. It will be understood

that these are merely illustrative external modules. Any other suitable modules may be used in addition to or in place of those illustrated.

Cipher feedback key generator 3014 may, externally to secure data parser 3000, generate for each secure data parser operation, a unique key, or random number (using, for example, random number generator 3012), to be used as a seed value for an operation that extends an original session key size (e.g., a value of 128, 256, 512, or 1024 bits) into a value equal to the length of the data to be parsed and split. Any suitable algorithm may be used for the cipher feedback key generation, including, for example, the AES cipher feedback key generation algorithm.

In order to facilitate integration of secure data parser 3000 and its external modules (i.e., secure data parser layer 3026) into an application layer 3024 (e.g., email application, database application, etc.), a wrapping layer that may make use of, for example, API function calls may be used. Any other suitable arrangement for facilitating integration of secure data parser layer 3026 into application layer 3024 may be used.

FIG. 31 illustratively shows how the arrangement of FIG. 30 may be used when a write (e.g., to a storage device), insert (e.g., in a database field), or transmit (e.g., across a network) command is issued in application layer 3024. At step 3100 data to be secured is identified and a call is made to the secure data parser. The call is passed through wrapper layer 3022 where at step 3102, wrapper layer 3022 streams the input data identified at step 3100 into assembled data buffer 3008. Also at step 3102, any suitable share information, filenames, any other suitable information, or any combination thereof may be stored (e.g., as information 3106 at wrapper layer 3022). Secure data processor 3000 then parses and splits the data it takes as input from assembled data buffer 3008 in accordance with the present invention. It outputs the data shares into split shares buffers 3010. At step 3104, wrapper layer 3022 obtains from stored information 3106 any suitable share information (i.e., stored by wrapper 3022 at step 3102) and share location(s) (e.g., from one or more configuration files). Wrapper layer 3022 then writes the output shares (obtained from split shares buffers 3010) appropriately (e.g., written to one or more storage devices, communicated onto a network, etc.).

FIG. 32 illustratively shows how the arrangement of FIG. 30 may be used when a read (e.g., from a storage device), select (e.g., from a database field), or receive (e.g., from a network) occurs. At step 3200, data to be restored is identified and a call to secure data parser 3000 is made from application layer 3024. At step 3202, from wrapper layer 3022, any suitable share information is obtained and share location is determined. Wrapper layer 3022 loads the portions of data identified at step 3200 into split shares buffers 3010. Secure data parser 3000 then processes these shares in accordance with the present invention (e.g., if only three of four shares are available, then the redundancy capabilities of secure data parser 3000 may be used to restore the original data using only the three shares). The restored data is then stored in assembled data buffer 3008. At step 3204, application layer 3022 converts the data stored in assembled data buffer 3008 into its original data format (if necessary) and provides the original data in its original format to application layer 3024.

It will be understood that the parsing and splitting of original data illustrated in FIG. 31 and the restoring of portions of data into original data illustrated in FIG. 32 is merely illustrative. Any other suitable processes, components, or both may be used in addition to or in place of those illustrated.

FIG. 33 is a block diagram of an illustrative process flow for parsing and splitting original data into two or more por-

tions of data in accordance with one embodiment of the present invention. As illustrated, the original data desired to be parsed and split is plain text **3306** (i.e., the word “SUMMIT” is used as an example). It will be understood that any other type of data may be parsed and split in accordance with the present invention. A session key **3300** is generated. If the length of session key **3300** is not compatible with the length of original data **3306**, then cipher feedback session key **3304** may be generated.

In one suitable approach, original data **3306** may be encrypted prior to parsing, splitting, or both. For example, as FIG. **33** illustrates, original data **3306** may be XORed with any suitable value (e.g., with cipher feedback session key **3304**, or with any other suitable value). It will be understood that any other suitable encryption technique may be used in place of or in addition to the XOR technique illustrate. It will further be understood that although FIG. **33** is illustrated in terms of byte by byte operations, the operation may take place at the bit level or at any other suitable level. It will further be understood that, if desired, there need not be any encryption whatsoever of original data **3306**.

The resultant encrypted data (or original data if no encryption took place) is then hashed to determine how to split the encrypted (or original) data among the output buckets (e.g., of which there are four in the illustrated example). In the illustrated example, the hashing takes place by bytes and is a function of cipher feedback session key **3304**. It will be understood that this is merely illustrative. The hashing may be performed at the bit level, if desired. The hashing may be a function of any other suitable value besides cipher feedback session key **3304**. In another suitable approach, hashing need not be used. Rather, any other suitable technique for splitting data may be employed.

FIG. **34** is a block diagram of an illustrative process flow for restoring original data **3306** from two or more parsed and split portions of original data **3306** in accordance with one embodiment of the present invention. The process involves hashing the portions in reverse (i.e., to the process of FIG. **33**) as a function of cipher feedback session key **3304** to restore the encrypted original data (or original data if there was no encryption prior to the parsing and splitting). The encryption key may then be used to restore the original data (i.e., in the illustrated example, cipher feedback session key **3304** is used to decrypt the XOR encryption by XORing it with the encrypted data). This restores original data **3306**.

FIG. **35** shows how bit-splitting may be implemented in the example of FIGS. **33** and **34**. A hash may be used (e.g., as a function of the cipher feedback session key, as a function of any other suitable value) to determine a bit value at which to split each byte of data. It will be understood that this is merely one illustrative way in which to implement splitting at the bit level. Any other suitable technique may be used.

It will be understood that any reference to hash functionality made herein may be made with respect to any suitable hash algorithm. These include for example, MD5 and SHA-1. Different hash algorithms may be used at different times and by different components of the present invention.

After a split point has been determined in accordance with the above illustrative procedure or through any other procedure or algorithm, a determination may be made with regard to which data portions to append each of the left and right segments. Any suitable algorithm may be used for making this determination. For example, in one suitable approach, a table of all possible distributions (e.g., in the form of pairings of destinations for the left segment and for the right segment) may be created, whereby a destination share value for each of the left and right segment may be determined by using any

suitable hash function on corresponding data in the session key, cipher feedback session key, or any other suitable random or pseudo-random value, which may be generated and extended to the size of the original data. For example, a hash function of a corresponding byte in the random or pseudo-random value may be made. The output of the hash function is used to determine which pairing of destinations (i.e., one for the left segment and one for the right segment) to select from the table of all the destination combinations. Based on this result, each segment of the split data unit is appended to the respective two shares indicated by the table value selected as a result of the hash function.

Redundancy information may be appended to the data portions in accordance with the present invention to allow for the restoration of the original data using fewer than all the data portions. For example, if two out of four portions are desired to be sufficient for restoration of data, then additional data from the shares may be accordingly appended to each share in, for example, a round-robin manner (e.g., where the size of the original data is 4 MB, then share 1 gets its own shares as well as those of shares 2 and 3; share 2 gets its own share as well as those of shares 3 and 4; share 3 gets its own share as well as those of shares 4 and 1; and share 4 gets its own shares as well as those of shares 1 and 2). Any such suitable redundancy may be used in accordance with the present invention.

It will be understood that any other suitable parsing and splitting approach may be used to generate portions of data from an original data set in accordance with the present invention. For example, parsing and splitting may be randomly or pseudo-randomly processed on a bit by bit basis. A random or pseudo-random value may be used (e.g., session key, cipher feedback session key, etc.) whereby for each bit in the original data, the result of a hash function on corresponding data in the random or pseudo-random value may indicate to which share to append the respective bit. In one suitable approach the random or pseudo-random value may be generated as, or extended to, 8 times the size of the original data so that the hash function may be performed on a corresponding byte of the random or pseudo-random value with respect to each bit of the original data. Any other suitable algorithm for parsing and splitting data on a bit by bit level may be used in accordance with the present invention. It will further be appreciated that redundancy data may be appended to the data shares such as, for example, in the manner described immediately above in accordance with the present invention.

In one suitable approach, parsing and splitting need not be random or pseudo-random. Rather, any suitable deterministic algorithm for parsing and splitting data may be used. For example, breaking up the original data into sequential shares may be employed as a parsing and splitting algorithm. Another example is to parse and split the original data bit by bit, appending each respective bit to the data shares sequentially in a round-robin manner. It will further be appreciated that redundancy data may be appended to the data shares such as, for example, in the manner described above in accordance with the present invention.

In one embodiment of the present invention, after the secure data parser generates a number of portions of original data, in order to restore the original data, certain one or more of the generated portions may be mandatory. For example, if one of the portions is used as an authentication share (e.g., saved on a physical token device), and if the fault tolerance feature of the secure data parser is being used (i.e., where fewer than all portions are necessary to restore the original data), then even though the secure data parser may have access to a sufficient number of portions of the original data in

order to restore the original data, it may require the authentication share stored on the physical token device before it restores the original data. It will be understood that any number and types of particular shares may be required based on, for example, application, type of data, user, any other suitable factors, or any combination thereof.

In one suitable approach, the secure data parser or some external component to the secure data parser may encrypt one or more portions of the original data. The encrypted portions may be required to be provided and decrypted in order to restore the original data. The different encrypted portions may be encrypted with different encryption keys. For example, this feature may be used to implement a more secure “two man rule” whereby a first user would need to have a particular share encrypted using a first encryption and a second user would need to have a particular share encrypted using a second encryption key. In order to access the original data, both users would need to have their respective encryption keys and provide their respective portions of the original data. In one suitable approach, a public key may be used to encrypt one or more data portions that may be a mandatory share required to restore the original data. A private key may then be used to decrypt the share in order to be used to restore to the original data.

Any such suitable paradigm may be used that makes use of mandatory shares where fewer than all shares are needed to restore original data.

In one suitable embodiment of the present invention, distribution of data into a finite number of shares of data may be processed randomly or pseudo-randomly such that from a statistical perspective, the probability that any particular share of data receives a particular unit of data is equal to the probability that any one of the remaining shares will receive the unit of data. As a result, each share of data will have an approximately equal amount of data bits.

According to another embodiment of the present invention, each of the finite number of shares of data need not have an equal probability of receiving units of data from the parsing and splitting of the original data. Rather certain one or more shares may have a higher or lower probability than the remaining shares. As a result, certain shares may be larger or smaller in terms of bit size relative to other shares. For example, in a two-share scenario, one share may have a 1% probability of receiving a unit of data whereas the second share has a 99% probability. It should follow, therefore that once the data units have been distributed by the secure data parser among the two share, the first share should have approximately 1% of the data and the second share 99%. Any suitable probabilities may be used in accordance with the present invention.

It will be understood that the secure data parser may be programmed to distribute data to shares according to an exact (or near exact) percentage as well. For example, the secure data parser may be programmed to distribute 80% of data to a first share and the remaining 20% of data to a second share.

According to another embodiment of the present invention, the secure data parser may generate data shares, one or more of which have predefined sizes. For example, the secure data parser may split original data into data portions where one of the portions is exactly 256 bits. In one suitable approach, if it is not possible to generate a data portion having the requisite size, then the secure data parser may pad the portion to make it the correct size. Any suitable size may be used.

In one suitable approach, the size of a data portion may be the size of an encryption key, a splitting key, any other suitable key, or any other suitable data element.

As previously discussed, the secure data parser may use keys in the parsing and splitting of data. For purposes of clarity and brevity, these keys shall be referred to herein as “splitting keys.” For example, the Session Master Key, previously introduced, is one type of splitting key. Also, as previously discussed, splitting keys may be secured within shares of data generated by the secure data parser. Any suitable algorithms for securing splitting keys may be used to secure them among the shares of data. For example, the Shamir algorithm may be used to secure the splitting keys whereby information that may be used to reconstruct a splitting key is generated and appended to the shares of data. Any other such suitable algorithm may be used in accordance with the present invention.

Similarly, any suitable encryption keys may be secured within one or more shares of data according to any suitable algorithm such as the Shamir algorithm. For example, encryption keys used to encrypt a data set prior to parsing and splitting, encryption keys used to encrypt a data portions after parsing and splitting, or both may be secured using, for example, the Shamir algorithm or any other suitable algorithm.

According to one embodiment of the present invention, an All or Nothing Transform (AoNT), such as a Full Package Transform, may be used to further secure data by transforming splitting keys, encryption keys, any other suitable data elements, or any combination thereof. For example, an encryption key used to encrypt a data set prior to parsing and splitting in accordance with the present invention may be transformed by an AoNT algorithm. The transformed encryption key may then be distributed among the data shares according to, for example, the Shamir algorithm or any other suitable algorithm. In order to reconstruct the encryption key, the encrypted data set must be restored (e.g., not necessarily using all the data shares if redundancy was used in accordance with the present invention) in order to access the necessary information regarding the transformation in accordance with AoNTs as is well known by one skilled in the art. When the original encryption key is retrieved, it may be used to decrypt the encrypted data set to retrieve the original data set. It will be understood that the fault tolerance features of the present invention may be used in conjunction with the AoNT feature. Namely, redundancy data may be included in the data portions such that fewer than all data portions are necessary to restore the encrypted data set.

It will be understood that the AoNT may be applied to encryption keys used to encrypt the data portions following parsing and splitting either in place of or in addition to the encryption and AoNT of the respective encryption key corresponding to the data set prior to parsing and splitting. Likewise, AoNT may be applied to splitting keys.

In one embodiment of the present invention, encryption keys, splitting keys, or both as used in accordance with the present invention may be further encrypted using, for example, a workgroup key in order to provide an extra level of security to a secured data set.

In one embodiment of the present invention, an audit module may be provided that tracks whenever the secure data parser is invoked to split data.

FIG. 36 illustrates possible options 3600 for using the components of the secure data parser in accordance with the invention. Each combination of options is outlined below and labeled with the appropriate step numbers from FIG. 36. The secure data parser may be modular in nature, allowing for any known algorithm to be used within each of the function blocks shown in FIG. 36. For example, other key splitting (e.g., secret sharing) algorithms such as Blakely may be used

in place of Shamir, or the AES encryption could be replaced by other known encryption algorithms such as Triple DES. The labels shown in the example of FIG. 36 merely depict one possible combination of algorithms for use in one embodiment of the invention. It should be understood that any suitable algorithm or combination of algorithms may be used in place of the labeled algorithms.

1) **3610, 3612, 3614, 3615, 3616, 3617, 3618, 3619**

Using previously encrypted data at step 3610, the data may be eventually split into a predefined number of shares. If the split algorithm requires a key, a split encryption key may be generated at step 3612 using a cryptographically secure pseudo-random number generator. The split encryption key may optionally be transformed using an All or Nothing Transform (AoNT) into a transform split key at step 3614 before being key split to the predefined number of shares with fault tolerance at step 3615. The data may then be split into the predefined number of shares at step 3616. A fault tolerant scheme may be used at step 3617 to allow for regeneration of the data from less than the total number of shares. Once the shares are created, authentication/integrity information may be embedded into the shares at step 3618. Each share may be optionally post-encrypted at step 3619.

2) **3111, 3612, 3614, 3615, 3616, 3617, 3618, 3619**

In some embodiments, the input data may be encrypted using an encryption key provided by a user or an external system. The external key is provided at step 3611. For example, the key may be provided from an external key store. If the split algorithm requires a key, the split encryption key may be generated using a cryptographically secure pseudo-random number generator at step 3612. The split key may optionally be transformed using an All or Nothing Transform (AoNT) into a transform split encryption key at step 3614 before being key split to the predefined number of shares with fault tolerance at step 3615. The data is then split to a predefined number of shares at step 3616. A fault tolerant scheme may be used at step 3617 to allow for regeneration of the data from less than the total number of shares. Once the shares are created, authentication/integrity information may be embedded into the shares at step 3618. Each share may be optionally post-encrypted at step 3619.

3) **3612, 3613, 3614, 3615, 3612, 3614, 3615, 3616, 3617, 3618, 3619**

In some embodiments, an encryption key may be generated using a cryptographically secure pseudo-random number generator at step 3612 to transform the data. Encryption of the data using the generated encryption key may occur at step 3613. The encryption key may optionally be transformed using an All or Nothing Transform (AoNT) into a transform encryption key at step 3614. The transform encryption key and/or generated encryption key may then be split into the predefined number of shares with fault tolerance at step 3615. If the split algorithm requires a key, generation of the split encryption key using a cryptographically secure pseudo-random number generator may occur at step 3612. The split key may optionally be transformed using an All or Nothing Transform (AoNT) into a transform split encryption key at step 3614 before being key split to the predefined number of shares with fault tolerance at step 3615. The data may then be split into a predefined number of shares at step 3616. A fault tolerant scheme may be used at step 3617 to allow for regeneration of the data from less than the total number of shares. Once the shares are created, authentication/integrity information will be embedded into the shares at step 3618. Each share may then be optionally post-encrypted at step 3619.

4) **3612, 3614, 3615, 3616, 3617, 3618, 3619**

In some embodiments, the data may be split into a predefined number of shares. If the split algorithm requires a key, generation of the split encryption key using a cryptographically secure pseudo-random number generator may occur at step 3612. The split key may optionally be transformed using an All or Nothing Transform (AoNT) into a transformed split key at step 3614 before being key split into the predefined number of shares with fault tolerance at step 3615. The data may then be split at step 3616. A fault tolerant scheme may be used at step 3617 to allow for regeneration of the data from less than the total number of shares. Once the shares are created, authentication/integrity information may be embedded into the shares at step 3618. Each share may be optionally post-encrypted at step 3619.

Although the above four combinations of options are preferably used in some embodiments of the invention, any other suitable combinations of features, steps, or options may be used with the secure data parser in other embodiments.

The secure data parser may offer flexible data protection by facilitating physical separation. Data may be first encrypted, then split into shares with "m of n" fault tolerance. This allows for regeneration of the original information when less than the total number of shares is available. For example, some shares may be lost or corrupted in transmission. The lost or corrupted shares may be recreated from fault tolerance or integrity information appended to the shares, as discussed in more detail below.

In order to create the shares, a number of keys are optionally utilized by the secure data parser. These keys may include one or more of the following:

Pre-encryption key: When pre-encryption of the shares is selected, an external key may be passed to the secure data parser. This key may be generated and stored externally in a key store (or other location) and may be used to optionally encrypt data prior to data splitting.

Split encryption key: This key may be generated internally and used by the secure data parser to encrypt the data prior to splitting. This key may then be stored securely within the shares using a key split algorithm.

Split session key: This key is not used with an encryption algorithm; rather, it may be used to key the data partitioning algorithms when random splitting is selected. When a random split is used, a split session key may be generated internally and used by the secure data parser to partition the data into shares. This key may be stored securely within the shares using a key splitting algorithm.

Post encryption key: When post encryption of the shares is selected, an external key may be passed to the secure data parser and used to post encrypt the individual shares. This key may be generated and stored externally in a key store or other suitable location.

In some embodiments, when data is secured using the secure data parser in this way, the information may only be reassembled provided that all of the required shares and external encryption keys are present.

FIG. 37 shows illustrative overview process 3700 for using the secure data parser of the present invention in some embodiments. As described above, two well-suited functions for secure data parser 3706 may include encryption 3702 and backup 3704. As such, secure data parser 3706 may be integrated with a RAID or backup system or a hardware or software encryption engine in some embodiments.

The primary key processes associated with secure data parser 3706 may include one or more of pre-encryption process 3708, encrypt/transform process 3710, key secure process 3712, parse/distribute process 3714, fault tolerance process

cess 3716, share authentication process 3716, and post-encryption process 3720. These processes may be executed in several suitable orders or combinations, as detailed in FIG. 36. The combination and order of processes used may depend on the particular application or use, the level of security desired, whether optional pre-encryption, post-encryption, or both, are desired, the redundancy desired, the capabilities or performance of an underlying or integrated system, or any other suitable factor or combination of factors.

The output of illustrative process 3700 may be two or more shares 3722. As described above, data may be distributed to each of these shares randomly (or pseudo-randomly) in some embodiments. In other embodiments, a deterministic algorithm (or some suitable combination of random, pseudo-random, and deterministic algorithms) may be used.

In addition to the individual protection of information assets, there is sometimes a requirement to share information among different groups of users or communities of interest. It may then be necessary to either control access to the individual shares within that group of users or to share credentials among those users that would only allow members of the group to reassemble the shares. To this end, a workgroup key may be deployed to group members in some embodiments of the invention. The workgroup key should be protected and kept confidential, as compromise of the workgroup key may potentially allow those outside the group to access information. Some systems and methods for workgroup key deployment and protection are discussed below.

The workgroup key concept allows for enhanced protection of information assets by encrypting key information stored within the shares. Once this operation is performed, even if all required shares and external keys are discovered, an attacker has no hope of recreating the information without access to the workgroup key.

FIG. 38 shows illustrative block diagram 3800 for storing key and data components within the shares. In the example of diagram 3800, the optional pre-encrypt and post-encrypt steps are omitted, although these steps may be included in other embodiments.

The simplified process to split the data includes encrypting the data using encryption key 3804 at encryption stage 3802. Portions of encryption key 3804 may then be split and stored within shares 3810 in accordance with the present invention. Portions of split encryption key 3806 may also be stored within shares 3810. Using the split encryption key, data 3808 is then split and stored in shares 3810.

In order to restore the data, split encryption key 3806 may be retrieved and restored in accordance with the present invention. The split operation may then be reversed to restore the ciphertext. Encryption key 3804 may also be retrieved and restored, and the ciphertext may then be decrypted using the encryption key.

When a workgroup key is utilized, the above process may be changed slightly to protect the encryption key with the workgroup key. The encryption key may then be encrypted with the workgroup key prior to being stored within the shares. The modified steps are shown in illustrative block diagram 3900 of FIG. 39.

The simplified process to split the data using a workgroup key includes first encrypting the data using the encryption key at stage 3902. The encryption key may then be encrypted with the workgroup key at stage 3904. The encryption key encrypted with the workgroup key may then be split into portions and stored with shares 3912. Split key 3908 may also be split and stored in shares 3912. Finally, portions of data 3910 are split and stored in shares 3912 using split key 3908.

In order to restore the data, the split key may be retrieved and restored in accordance with the present invention. The split operation may then be reversed to restore the ciphertext in accordance with the present invention. The encryption key (which was encrypted with the workgroup key) may be retrieved and restored. The encryption key may then be decrypted using the workgroup key. Finally, the ciphertext may be decrypted using the encryption key.

There are several secure methods for deploying and protecting workgroup keys. The selection of which method to use for a particular application depends on a number of factors. These factors may include security level required, cost, convenience, and the number of users in the workgroup. Some commonly used techniques used in some embodiments are provided below:

Hardware-Based Key Storage

Hardware-based solutions generally provide the strongest guarantees for the security of encryption/decryption keys in an encryption system. Examples of hardware-based storage solutions include tamper-resistant key token devices which store keys in a portable device (e.g., smartcard/dongle), or non-portable key storage peripherals. These devices are designed to prevent easy duplication of key material by unauthorized parties. Keys may be generated by a trusted authority and distributed to users, or generated within the hardware. Additionally, many key storage systems provide for multi-factor authentication, where use of the keys requires access both a physical object (token) and a passphrase or biometric.

Software-Based Key Storage

While dedicated hardware-based storage may be desirable for high-security deployments or applications, other deployments may elect to store keys directly on local hardware (e.g., disks, RAM or non-volatile RAM stores such as USB drives). This provides a lower level of protection against insider attacks, or in instances where an attacker is able to directly access the encryption machine.

To secure keys on disk, software-based key management often protects keys by storing them in encrypted form under a key derived from a combination of other authentication metrics, including: passwords and passphrases, presence of other keys (e.g., from a hardware-based solution), biometrics, or any suitable combination of the foregoing. The level of security provided by such techniques may range from the relatively weak key protection mechanisms provided by some operating systems (e.g., MS Windows and Linux), to more robust solutions implemented using multi-factor authentication.

The secure data parser of the present invention may be advantageously used in a number of applications and technologies. For example, email system, RAID systems, video broadcasting systems, database systems, tape backup systems, or any other suitable system may have the secure data parser integrated at any suitable level. As previously discussed, it will be understood that the secure data parser may also be integrated for protection and fault tolerance of any type of data in motion through any transport medium, including, for example, wired, wireless, or physical transport mediums. As one example, voice over Internet protocol (VoIP) applications may make use of the secure data parser of the present invention to solve problems relating to echoes and delays that are commonly found in VoIP. The need for network retry on dropped packets may be eliminated by using fault tolerance, which guarantees packet delivery even with the loss of a predetermined number of shares. Packets of data (e.g., network packets) may also be efficiently split and restored "on-the-fly" with minimal delay and buffering, resulting in a comprehensive solution for various types of data

in motion. The secure data parser may act on network data packets, network voice packets, file system data blocks, or any other suitable unit of information. In addition to being integrated with a VoIP application, the secure data parser may be integrated with a file-sharing application (e.g., a peer-to-peer file-sharing application), a video broadcasting application, an electronic voting or polling application (which may implement an electronic voting protocol and blind signatures, such as the Sensus protocol), an email application, or any other network application that may require or desire secure communication.

In some embodiments, support for network data in motion may be provided by the secure data parser of the present invention in two distinct phases—a header generation phase and a data partitioning phase. Simplified header generation process **4000** and simplified data partitioning process **4010** are shown in FIGS. **40A** and **40B**, respectively. One or both of these processes may be performed on network packets, file system blocks, or any other suitable information.

In some embodiments, header generation process **4000** may be performed one time at the initiation of a network packet stream. At step **4002**, a random (or pseudo-random) split encryption key, *K*, may be generated. The split encryption key, *K*, may then be optionally encrypted (e.g., using the workgroup key described above) at AES key wrap step **4004**. Although an AES key wrap may be used in some embodiments, any suitable key encryption or key wrap algorithm may be used in other embodiments. AES key wrap step **4004** may operate on the entire split encryption key, *K*, or the split encryption key may be parsed into several blocks (e.g., 64-bit blocks). AES key wrap step **4004** may then operate on blocks of the split encryption key, if desired.

At step **4006**, a secret sharing algorithm (e.g., Shamir) may be used to split the split encryption key, *K*, into key shares. Each key share may then be embedded into one of the output shares (e.g., in the share headers). Finally, a share integrity block and (optionally) a post-authentication tag (e.g., MAC) may be appended to the header block of each share. Each header block may be designed to fit within a single data packet.

After header generation is complete (e.g., using simplified header generation process **4000**), the secure data parser may enter the data partitioning phase using simplified data splitting process **4010**. Each incoming data packet or data block in the stream is encrypted using the split encryption key, *K*, at step **4012**. At step **4014**, share integrity information (e.g., a hash *H*) may be computed on the resulting ciphertext from step **4012**. For example, a SHA-256 hash may be computed. At step **4106**, the data packet or data block may then be partitioned into two or more data shares using one of the data splitting algorithms described above in accordance with the present invention. In some embodiments, the data packet or data block may be split so that each data share contains a substantially random distribution of the encrypted data packet or data block. The integrity information (e.g., hash *H*) may then be appended to each data share. An optional post-authentication tag (e.g., MAC) may also be computed and appended to each data share in some embodiments.

Each data share may include metadata, which may be necessary to permit correct reconstruction of the data blocks or data packets. This information may be included in the share header. The metadata may include such information as cryptographic key shares, key identities, share nonces, signatures/MAC values, and integrity blocks. In order to maximize bandwidth efficiency, the metadata may be stored in a compact binary format.

For example, in some embodiments, the share header includes a cleartext header chunk, which is not encrypted and may include such elements as the Shamir key share, per-session nonce, per-share nonce, key identifiers (e.g., a workgroup key identifier and a post-authentication key identifier). The share header may also include an encrypted header chunk, which is encrypted with the split encryption key. An integrity header chunk, which may include integrity checks for any number of the previous blocks (e.g., the previous two blocks) may also be included in the header. Any other suitable values or information may also be included in the share header.

As shown in illustrative share format **4100** of FIG. **41**, header block **4102** may be associated with two or more output blocks **4104**. Each header block, such as header block **4102**, may be designed to fit within a single network data packet. In some embodiments, after header block **4102** is transmitted from a first location to a second location, the output blocks may then be transmitted. Alternatively, header block **4102** and output blocks **4104** may be transmitted at the same time in parallel. The transmission may occur over one or more similar or dissimilar communications paths.

Each output block may include data portion **4106** and integrity/authenticity portion **4108**. As described above, each data share may be secured using a share integrity portion including share integrity information (e.g., a SHA-256 hash) of the encrypted, pre-partitioned data. To verify the integrity of the outputs blocks at recovery time, the secure data parser may compare the share integrity blocks of each share and then invert the split algorithm. The hash of the recovered data may then be verified against the share hash.

As previously mentioned, in some embodiments of the present invention, the secure data parser may be used in conjunction with a tape backup system. For example, an individual tape may be used as a node (i.e., portion/share) in accordance with the present invention. Any other suitable arrangement may be used. For example, a tape library or subsystem, which is made up of two or more tapes, may be treated as a single node.

Redundancy may also be used with the tapes in accordance with the present invention. For example, if a data set is apportioned among four tapes (i.e., portions/shares), then two of the four tapes may be necessary in order to restore the original data. It will be understood that any suitable number of nodes (i.e., less than the total number of nodes) may be required to restore the original data in accordance with the redundancy features of the present invention. This substantially increases the probability for restoration when one or more tapes expire.

Each tape may also be digitally protected with a SHA-256, HMAC hash value, any other suitable value, or any combination thereof to insure against tampering. Should any data on the tape or the hash value change, that tape would not be a candidate for restoration and any minimum required number of tapes of the remaining tapes would be used to restore the data.

In conventional tape backup systems, when a user calls for data to be written to or read from a tape, the tape management system (TMS) presents a number that corresponds to a physical tape mount. This tape mount points to a physical drive where the data will be mounted. The tape is loaded either by a human tape operator or by a tape robot in a tape silo.

Under the present invention, the physical tape mount may be considered a logical mount point that points to a number of physical tapes. This not only increases the data capacity but also improves the performance because of the parallelism.

For increased performance the tape nodes may be or may include a RAID array of disks used for storing tape images.

This allows for high-speed restoration because the data may always be available in the protected RAID.

In any of the foregoing embodiments, the data to be secured may be distributed into a plurality of shares using deterministic, probabilistic, or both deterministic and probabilistic data distribution techniques. In order to prevent an attacker from beginning a crypto attack on any cipher block, the bits from cipher blocks may be deterministically distributed to the shares. For example, the distribution may be performed using the BitSegment routine, or the BlockSegment routine may be modified to allow for distribution of portions of blocks to multiple shares. This strategy may defend against an attacker who has accumulated less than “M” shares.

In some embodiments, a keyed secret sharing routine may be employed using keyed information dispersal (e.g., through the use of a keyed information dispersal algorithm or “IDA”). The key for the keyed IDA may also be protected by one or more external workgroup keys, one or more shared keys, or any combination of workgroup keys and shared keys. In this way, a multi-factor secret sharing scheme may be employed. To reconstruct the data, at least “M” shares plus the workgroup key(s) (and/or shared key(s)) may be required in some embodiments. The IDA (or the key for the IDA) may also be driven into the encryption process. For example, the transform may be driven into the clear text (e.g., during the pre-processing layer before encrypting) and may further protect the clear text before it is encrypted.

For example, in some embodiments, keyed information dispersal is used to distribute unique portions of data from a data set into two or more shares. The keyed information dispersal may use a session key to first encrypt the data set, to distribute unique portions of encrypted data from the data set into two or more encrypted data set shares, or both encrypt the data set and distribute unique portions of encrypted data from the data set into the two or more encrypted data set shares. For example, to distribute unique portions of the data set or encrypted data set, secret sharing (or the methods described above, such as BitSegment or BlockSegment) may be used. The session key may then optionally be transformed (for example, using a full package transform or AoNT) and shared using, for example, secret sharing (or the keyed information dispersal and session key).

In some embodiments, the session key may be encrypted using a shared key (e.g., a workgroup key) before unique portions of the key are distributed or shared into two or more session key shares. Two or more user shares may then be formed by combining at least one encrypted data set share and at least one session key share. In forming a user share, in some embodiments, the at least one session key share may be interleaved into an encrypted data set share. In other embodiments, the at least one session key share may be inserted into an encrypted data set share at a location based at least in part on the shared workgroup key. For example, keyed information dispersal may be used to distribute each session key share into a unique encrypted data set share to form a user share. Interleaving or inserting a session key share into an encrypted data set share at a location based at least in part on the shared workgroup may provide increased security in the face of cryptographic attacks. In other embodiments, one or more session key shares may be appended to the beginning or end of an encrypted data set share to form a user share. The collection of user shares may then be stored separately on at least one data depository. The data depository or depositories may be located in the same physical location (for example, on the same magnetic or tape storage device) or geographically separated (for example, on physically separated servers in different geographic locations). To reconstruct the original

data set, an authorized set of user shares and the shared workgroup key may be required.

Keyed information dispersal may be secure even in the face of key-retrieval oracles. For example, take a blockcipher E and a key-retrieval oracle for E that takes a list $(X_1, Y_1), \dots, (X_c, Y_c)$ of input/output pairs to the blockcipher, and returns a key K that is consistent with the input/output examples (e.g., $Y_i = E_K(X_i)$ for all i). The oracle may return the distinguished value \perp if there is no consistent key. This oracle may model a cryptanalytic attack that may recover a key from a list of input/output examples.

Standard blockcipher-based schemes may fail in the presence of a key-retrieval oracle. For example, CBC encryption or the CBC MAC may become completely insecure in the presence of a key-retrieval oracle.

If Π^{IDA} is an IDA scheme and Π^{Enc} is an encryption scheme given by a mode of operation of some blockcipher E, then (Π^{IDA}, Π^{Enc}) provides security in the face of a key-retrieval attack if the two schemes, when combined with an arbitrary perfect secret-sharing scheme (PSS) as per HK1 or HK2, achieve the robust computational secret sharing (RCSS) goal, but in the model in which the adversary has a key-retrieval oracle.

If there exists an IDA scheme Π^{IDA} and an encryption scheme Π^{Enc} such that the pair of schemes provides security in the face of key-retrieval attacks, then one way to achieve this pair may be to have a “clever” IDA and a “dumb” encryption scheme. Another way to achieve this pair of schemes may be to have a “dumb” IDA and a “clever” encryption scheme.

To illustrate the use of a clever IDA and a dumb encryption scheme, in some embodiments, the encryption scheme may be CBC and the IDA may have a “weak privacy” property. The weak privacy property means, for example, that if the input to the IDA is a random sequence of blocks $M = M_1 \dots M_l$ and the adversary obtains shares from a non-authorized collection, then there is some block index i such that it is infeasible for the adversary to compute M_i . Such a weakly-private IDA may be built by first applying to Man information-theoretic AoNT, such as Stinson’s AoNT, and then applying a simple IDA such as BlockSegment, or a bit-efficient IDA like Rabin’s scheme (e.g., Reed-Solomon encoding).

To illustrate the use of a dumb IDA and a clever encryption scheme, in some embodiments, one may use a CBC mode with double encryption instead of single encryption. Now any IDA may be used, even replication. Having the key-retrieval oracle for the blockcipher would be useless to an adversary, as the adversary will be denied any singly-enciphered input/output example.

While a clever IDA has value, it may also be inessential in some contexts, in the sense that the “smarts” needed to provide security in the face of a key-retrieval attack could have been “pushed” elsewhere. For example, in some embodiments, no matter how smart the IDA, and for whatever goal is trying to be achieved with the IDA in the context of HK1/HK2, the smarts may be pushed out of the IDA and into the encryption scheme, being left with a fixed and dumb IDA.

Based on the above, in some embodiments, a “universally sound” clever IDA Π^{IDA} may be used. For example, an IDA is provided such that, for all encryption schemes Π^{Enc} , the pair (Π^{IDA}, Π^{Enc}) universally provides security in the face of key-retrieval attacks.

In some embodiments, an encryption scheme is provided that is RCSS secure in the face of a key-retrieval oracle. The scheme may be integrated with HK1/HK2, with any IDA, to achieve security in the face of key-retrieval. Using the new

scheme may be particularly useful, for example, for making symmetric encryption schemes more secure against key-retrieval attacks.

As mentioned above, classical secret-sharing notions are typically unkeyed. Thus, a secret is broken into shares, or reconstructed from them, in a way that requires neither the dealer nor the party reconstructing the secret to hold any kind of symmetric or asymmetric key. The secure data parser described herein, however, is optionally keyed. The dealer may provide a symmetric key that, if used for data sharing, may be required for data recovery. The secure data parser may use the symmetric key to disperse or distribute unique portions of the message to be secured into two or more shares.

The shared key may enable multi-factor or two-factor secret-sharing (2FSS). The adversary may then be required to navigate through two fundamentally different types of security in order to break the security mechanism. For example, to violate the secret-sharing goals, the adversary (1) may need to obtain the shares of an authorized set of players, and (2) may need to obtain a secret key that it should not be able to obtain (or break the cryptographic mechanism that is keyed by that key).

In some embodiments, a new set of additional requirements is added to the RCSS goal. The additional requirements may include the “second factor”—key possession. These additional requirements may be added without diminishing the original set of requirements. One set of requirements may relate to the adversary’s inability to break the scheme if it knows the secret key but does not obtain enough shares (e.g., the classical or first factor requirements) while the other set of requirements may relate to the adversary’s inability to break the scheme if it does have the secret key but manages to get hold of all of the shares (e.g., the new or second-factor requirements).

In some embodiments, there may be two second-factor requirements: a privacy requirement and an authenticity requirement. In the privacy requirement, a game may be involved where a secret key K and a bit b are selected by the environment. The adversary now supplies a pair of equal-length messages in the domain of the secret-sharing scheme, M_1^0 and M_1^1 . The environment computes the shares of M_1^b to get a vector of shares, $S_1 = (S_1[1], \dots, S_1[n])$, and it gives the shares S_1 (all of them) to the adversary. The adversary may now choose another pair of messages (M_2^0, M_2^1) and everything proceeds as before, using the same key K and hidden bit b . The adversary’s job is to output the bit b' that it believes to be b . The adversary privacy advantage is one less than twice the probability that $b=b'$. This game captures the notion that, even learning all the shares, the adversary still cannot learn anything about the shared secret if it lacks the secret key.

In the authenticity requirement, a game may be involved where the environment chooses a secret key K and uses this in the subsequent calls to Share and Recover. Share and Recover may have their syntax modified, in some embodiments, to reflect the presence of this key. Then the adversary makes Share requests for whatever messages M_1, \dots, M_q it chooses in the domain of the secret-sharing scheme. In response to each Share request it gets the corresponding n -vector of shares, S_1, \dots, S_q . The adversary’s aim is to forge a new plaintext; it wins if it outputs a vector of shares S' such that, when fed to the Recover algorithm, results in something not in $\{M_1, \dots, M_q\}$. This is an “integrity of plaintext” notion.

There are two approaches to achieve multi-factor secret-sharing. The first is a generic approach—generic in the sense of using an underlying (R)CSS scheme in a black-box way. An authenticated-encryption scheme is used to encrypt the message that is to be CSS-shared, and then the resulting

ciphertext may be shared out, for example, using a secret sharing algorithm, such as Blakely or Shamir.

A potentially more efficient approach is to allow the shared key to be the workgroup key. Namely, (1) the randomly generated session key of the (R)CSS scheme may be encrypted using the shared key, and (2) the encryption scheme applied to the message (e.g., the file) may be replaced by an authenticated-encryption scheme. This approach may entail only a minimal degradation in performance.

Although some applications of the secure data parser are described above, it should be clearly understood that the present invention may be integrated with any network application in order to increase security, fault-tolerance, anonymity, or any suitable combination of the foregoing.

In some embodiments of the present invention, the secure data parser may be implemented in a secure proxy service to secure data in motion. As described above, the secure data parser is a cryptographic library that provides traditional encryption/authentication services for applications, as well as an additional security property achieved by separating protected data (either physically, temporally, or by some other form of trust). The secure data parser is designed for applications where the threat of an adversary compromising the system is real, either by obtaining cryptographic keys, physical access to a transmission medium, or obtaining any knowledge that would ordinarily defeat the security. The secure proxy service provides an additional layer of security to protect from these same threats, and is preferably flexible such that it can be implemented on a wide range of systems—(e.g., enterprise servers, personal computers, any other suitable system, or any combination thereof) The secure proxy service is described with respect to FIGS. 42-50 below.

The secure proxy service is used to secure data in motion between two devices. In particular, the secure proxy service runs on a first device and provides secure data parser-enabled communications for applications over a network. These devices may be any suitable pair of devices included in cryptographic system 100 (FIG. 1). For example the secure proxy service may be established between user system 105 and vendor system 120, such as a personal computer and a web server. In another example, the secure proxy service may be established between separate user systems 105, such as a personal computer and a NAS, a personal computer and a home router, a NAS and a home router, or any suitable combination of user systems 105. The communication between devices using the secure proxy service resembles that of a client connecting to a web or e-mail server.

In some embodiments, a client, such as user system 105, and a server, such as vendor system 120, may establish secure communications using the secure proxy service. In establishing the secure proxy service, the user system 105 and the vendor system 120 may be retrofit to a suitable configuration for the secure proxy service. In such embodiments, when the client connects to the server, the connection is established between the two secure proxy services. On the server, the secure proxy service is configured to forward data it receives to a server application. The server application may then handle the request and respond through the locally implemented secure proxy service.

In some embodiments, the secure proxy service protects the confidentiality, integrity, and authenticity of the data transmitted over a network based on distributed trust among any number of certificate authorities, such as certificate authorities 115 (FIG. 1). In such embodiments, the confidentiality, integrity, and authenticity of the data may be protected so long as a quorum of certificate authorities is trusted. If the trust of the certificate authority is compromised, mutual

authentication cannot be assured, and the confidentiality, integrity, and authenticity of exchanged messages breaks down. The secure proxy service is a secure data parser enabled solution that allows the trust placed in a single certificate authority to be distributed over any number of certificate authorities (e.g., two, three, five, ten, twenty, fifty, one hundred, or more than one hundred certificate authorities). This distributed trust allows the exchange of messages over the secure proxy service to remain secure if there is a single point of failure among the set of certificate authorities.

In some embodiments, the secure proxy service is implemented using an adaptation of SSL and/or full TLS protocols. These protocols may be suitable for adaptation as part of the secure proxy service because they rely at least in part on the trust of a certificate authority for mutually authenticating both parties in a communication.

An overview of the use of certificate authorities in full TLS is now described. For a full TLS-enabled connection to be established between two devices, the two devices agree on the cryptographic suite of algorithms to use, and exchange and mutually authenticate one another's public keys. The public keys of each device are authenticated through validation of a certificate authority's signature of the public key. Trust that the two devices are genuinely communicating with one another is established by the fact that both of them trust the certificate authority, whose signature of the devices' certificates could not be forged without the compromise of that authority.

The certificate authority creates for itself a public and private key pair, (Pub_{CA}, Pri_{CA}). In addition, the certificate authority creates a self signed certificate for the public key:

$$\text{Cert}_{CA} = \text{Pub}_{CA}, \text{Sig}_{\text{Pri}_{CA}}(\text{Pub}_{CA}) \quad (1)$$

Both devices receive the certificate authority's certificate Cert_{CA} according to equation (1), a private key (Pri_{Dev1}, Pri_{Dev2}), and a certificate signed by the certificate authority:

$$\text{Cert}_{Dev1} = \text{Pub}_{Dev1}, \text{Sig}_{\text{Pri}_{CA}}(\text{Pub}_{Dev1}) \quad (2)$$

$$\text{Cert}_{Dev2} = \text{Pub}_{Dev2}, \text{Sig}_{\text{Pri}_{CA}}(\text{Pub}_{Dev2}) \quad (3)$$

When a communication begins, the devices exchange their respective certificates in equations (2) and (3), and verify the authenticity of these certificates using the public key of the certificate authority. For example, the first device may perform the verification by running a verification function Verify (Cert_{Dev2}, Pub_{CA}), and the second device may perform the verification by running a verification function Verify (Cert_{Dev1}, Pub_{CA}). If both devices are satisfied with the certificate authority's signature of the exchanged public keys, the first device sends the second device symmetric encryption key material using the second device's public key. The first device proves knowledge of the private key corresponding to their certificate by performing a digital signature challenge. Once the first device proves knowledge of the private key corresponding to their certificate, the first device and the second device may exchange messages securely.

If either the first device or the second device has been compromised, their respective private keys may be compromised as well and from then on the compromised device could be impersonated. If the certificate authority is compromised, valid certificates may be generated for which the certificate authority knows the corresponding private key, and either of the devices could be impersonated. However, absent compromise of the devices or the certificate authority, the devices can be mutually assured that they are talking to the correct entity.

In some embodiments, not all trust is delegated to a single certificate authority. It may be impractical to constantly create, distribute, and validate certificates by a single certificate authority. Instead, a chain of trust may be established in the form of a certificate authority hierarchy 4200, as illustrated by FIG. 42. Certificate authority hierarchy 4200 may establish a chain of trust in the form of a tree. At the top of the tree is root certificate authority 4210 who delegates authority to all descendants (e.g., children and grandchildren) of the root certificate authority 4210. The trust at each level of certificate authority hierarchy 4200 is assured by the trust of the root certificate authority 4210.

In certificate authority hierarchy 4200, root certificate authority 4210 may sign certificates for children certificate authorities 4220. Although only one root certificate authority CA-0 is shown in FIG. 42, it may be understood that in certain embodiments the certificate authority hierarchy may include any number of root certificate authorities. Children certificate authorities 4220 may sign certificates for grandchildren certificate authorities 4230. For example, as illustrated in FIG. 42, root certificate authority CA-0 signs certificates for its children CA-1 and CA-2, who in turn sign certificates for their children CA-1.1, CA-1.2, CA-2.1, and CA-2.2. Although only three levels of certificate authorities are illustrated in FIG. 42, it will be understood that in certain embodiments there may be greater or fewer levels of certificate authorities. To ensure nonrepudiation of signatures, all entities may generate their own certificates.

In a cryptographic system, a first device may have received its certificate from one of the children certificate authorities 4220, and a second device in communication with the first device may receive its certificate from one of the grandchildren certificate authorities 4230. Validation of the certificates for each device may be performed by obtaining the certificate of the issuing certificate authority (e.g., CA-1 or CA-2 for a first device, and CA-1.1, CA-1.2, CA-2.1, or CA-2.2 for a second device) and verifying the signature of the first device or second device's certificate (e.g., verifying the certificates shown in equations (2) or (3)). If the trust of the issuing certificate authority cannot be established, the device performing verification can obtain the certificate of the parent of the issuing certificate authority in question and perform a similar verification to ensure that the certificate authority is valid. This process may continue by both devices until reaching root certificate authority 4210, which is trusted by both devices. In some embodiments, each device that is in communication may be associated with more than one root certificate authorities 4210. In such embodiments, it is possible for devices with valid certificates from any of these certificate authorities to communicate.

From the above description of the use of certificate authorities with TLS, it is understood that security ultimately lays in the trust of a single root certificate authority, or to a lesser extent, one of the descendant certificate authorities within the hierarchy of certificate authorities that the devices in communication are comfortable trusting. In some embodiments, if any certificate authority in the hierarchy is compromised, all descendants of the compromised certificate authority are also compromised. If this compromised node is the root, then each of the certificate authorities in the hierarchy may be compromised.

In some embodiments, the secure proxy service may use the secure data parser with TLS to distribute the trust placed in a single certificate authority with the trust of a quorum of certificate authorities. This quorum may be a quorum of root certificate authorities 4210, or a quorum of minor certificate authorities within the tree of a single root certificate authority.

For example, this quorum may be two out of the three certificate authorities in the set consisting of CA-1, CA-1.1 and CA-1.2, which are minor certificate authorities within the tree of root certificate authority CA-0.

In some embodiments, certificate authority hierarchy **4200** may be traversed by any suitable graph algorithm. This traversal may be performed in order to obtain a list of certificate authorities or a list of certificates associated with certificate authorities that unique, or have different public and private key pairs. In some embodiments, the traversal of certificate authority hierarchy **4200** may result in certificate authorities or certificates of certificate authorities that are root certificate authorities. In some embodiments, the traversal of certificate authority hierarchy **4200** may result in certificate authorities or certificates of certificate authorities that are minor certificate authorities within the tree of one or more root certificate authorities.

FIGS. **44** through **50** detail two approaches to implementing the secure proxy service. Both approaches are equally secure. In some embodiments, the secure data parser may be integrated with full TLS. In addition, in each approach trust is distributed among a set of certificate authorities (e.g., the quorum of certificate authorities discussed with respect to certificate authority hierarchy **4200** in FIG. **42**). In some embodiments, the secure proxy service may be implemented by integrating the secure data parser with SSL, with SSL and TLS, or implementing the secure data parser without the use of SSL and/or TLS. In some embodiments, the secure proxy service may be implemented in conjunction with any one or more types of encryption **3018** that may provide secure encryption of data at the secure data parser layer **3026** of FIG. **30**. In addition, in some embodiments the secure proxy service may be implemented in conjunction with any suitable protocol that makes use of certificate authorities to ensure the confidentiality, integrity, and authenticity of exchanged messages.

In the embodiments described with respect to FIGS. **44** through **46**, the secure data parser may be used to distribute trust in any number of certificate authorities during initial negotiation (e.g., the key establishment phase) of a connection between devices. This offers the assurance that if some (but fewer than a quorum) of the certificate authorities have been compromised, the connection may still be established, and messages may be exchanged without disrupting the confidentiality, integrity, and authenticity of the communication. In the embodiments described with respect to FIGS. **47** through **50**, the data is pre-processed using the secure data parser and then dispersed into shares. A set of secure communication tunnels may be established within a communication channel using certificates issued by unique certificate authorities, these certificate authorities may be used to encrypt data for each of the tunnels, and the individual shares of data may be transmitted on each of the tunnels. Thus, in the second approach trust may be distributed among a set of certificate authorities in the structure of the communication channel itself.

In order to illustrate how the secure data parser is integrated with TLS in some embodiments of the secure proxy service, an overview of the key establishment phase of full TLS is described with respect to FIG. **43**. FIG. **43** shows a simplified and illustrative process flow **4300** for the key establishment phase of TLS for communication between two devices: first device **4310** and second device **4320**. First device **4310** and second device **4320** may be any combination of user system **105** and/or vendor system **120** communicating over a communication link, for example, communication link **125** as shown in FIG. **1**. This key establishment phase may include a

handshake and mutual authentication. At step **4312**, first device **4310** generates a random number R_{Dev1} and sends the random number along with its certificate $Cert_{Dev1}$ (as calculated, for example, in equation (2)) to second device **4320**.

At step **4322**, second device **4320** generates its own random number R_{Dev2} , and sends the random number along with its certificate $Cert_{Dev2}$ (as calculated in equation (3)). At step **4314**, the client generates secret information Sp_{Dev1} , encrypts it under the second device's public key using any suitable type of encryption, and sends it to the second device. At step **4324**, second device **4320** decrypts the secret information S_{Dev1} and computes a shared encryption key K based on a pseudo random function G and the random and secret values that have been exchanged (i.e., R_{Dev1} , R_{Dev2} , and S_{Dev1}). Similarly, at step **4316** second device **4310** computes a shared encryption key K based on a pseudo random function G and the random and secret values that have been exchanged (i.e., R_{Dev1} , R_{Dev2} , and S_{Dev1}). At step **4330**, first device **4310** and second device **4320** exchange messages encrypted with their independently calculated shared encryption keys. If the computed shared encryption keys match, first device **4310** and second device **4320** may exchange messages that are ensured to be confidential and authentic. As will be discussed with respect to FIGS. **44** through **46**, in some embodiments, the secure data parser service may modify and/or add to the steps of process flow **4300** in order to integrate the secure data parser with TLS.

FIG. **44** shows a simplified and illustrative process flow **4400** for a secure proxy service, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. In process flow **4400**, trust is distributed in a set of certificate authorities during initial negotiation of a connection between devices. In some embodiments, process flow **4400** may be executed as part of the key establishment phase of a secure exchange of information between two devices. This key establishment phase may be part of one or more of the processes associated with secure data parser **3706** as illustrated in FIG. **37**, or may be a standalone process. For example, steps **4410**, **4420**, **4430**, **4440**, **4450**, and **4460** may be part of one or more of pre-encryption process **3708**, encrypt/transform process **3710**, key secure process **3712**, or parse/distribute process **3714** associated with secure data parser **3706** as illustrated in FIG. **37**, or may be a standalone process.

Process flow **4400** begins at step **4410**. At step **4410**, a first device that would like to securely exchange information with device second device may generate secret information. This secret information may be any amount of suitable random numbers (e.g., one, two, five, twenty, one-hundred, or more than one-hundred random numbers) generated by a random number generator. For example, the secret information may be a random number generated by the random number generator **3012** of secure data parser **3026** as shown in FIG. **30**. Process flow **4400** proceeds to step **4420**.

At step **4420**, the first device may disperse the secret information generated at step **4410** into shares. In some embodiments, the secret information may be dispersed into shares using a cryptosplitting process, such as an "M of N cryptosplit". This "M of N cryptosplit" may be achieved using the secure data parser of the present invention. For example, the cryptosplit may be achieved using any of the data splitting techniques discussed with respect to FIG. **21** through FIG. **24**. In such embodiments, the dispersed shares may be restorable from at least a subset of the shares by recombining at least a quorum of the shares. In addition, in some embodiments, the split of secret information may occur

substantially through any number of uses of the secure data parser outlined with respect to FIG. 33, FIG. 35, and FIG. 36. For example, the secure data parser may receive unencrypted secret information at step 3610. If the secret information is going to be split with an algorithm that requires a key, a split encryption key is generated at step 3612. The secret information may be split into shares at step 3616 (e.g., according to any of the techniques described with respect to FIG. 33, FIG. 35, and FIG. 36). A fault tolerant scheme may be used at step 3617 to encrypt the split encryption key and allow for regeneration of the secret information from less than the total number of shares. In addition, at step 3617 information may be added to the shares of secret information that is used to reconstruct the shares. In some embodiments, this information may be embedded into share headers. Further, once the shares are created, authentication/integrity information may be embedded into the headers of the shares of secret information at step 3618. Each share may be post-encrypted using public keys of different certificate authorities as will be described with respect to step 4430.

In addition, in some embodiments, the dispersing of the secret information into shares may occur, for example, according to the simplified header generation process 4000 as shown in FIG. 40A. For example, at step 4002, the secret information may be generated. The secret information may then be optionally encrypted (e.g., using the workgroup key described with respect to FIG. 39) at step 4004. At step 4006, a "M of N cryptosplit" may be used to split the secret information into shares of secret information. Information associated with the split of secret information may then be embedded into a share header. Finally, a share integrity block and a post-authentication tag (e.g., MAC) may be appended to the header block of each share. Each header block may be designed to fit within a single data packet.

In some embodiments, the shares of secret information generated at step 4420 may be generated using a multi-factor secret sharing scheme. This multi-factor secret sharing scheme may be, for example, the keyed Information Dispersal Algorithm discussed after FIG. 41. For example, the shares of secret information may be distributed with the data to be secured into a plurality of shares using deterministic, probabilistic, or both deterministic and probabilistic data distribution techniques. Once the secret information has been dispersed into shares, process flow 4400 may proceed to step 4430

At step 4430, the shares resulting from the split of secret information computed in step 4420 are encrypted by the first device based on public keys of unique certificate authorities. For example, if there were three shares of secret information, the first share may be encrypted under the public key of a first certificate authority, the second share may be encrypted under the public key of a second certificate authority, and the third share may be encrypted under the public key of a third certificate authority. Each certificate authority may be unique in that the certificates issued by each have different public and private key pairs. In some embodiments, the unique certificate authorities may be root certificate authorities. In other embodiments, the unique certificate authorities may be minor certificate authorities within the tree of a single root certificate authority, as discussed with respect to certificate authority hierarchy 4200 in FIG. 42.

As discussed above with respect to step 4420, in some embodiments information related to the dispersal of secret information may be embedded into share headers. For example, if the secret information is split at step 4420 into

four shares, four headers may be generated that each includes information associated with the dispersed shares of secret information.

In some embodiments, the shares may be protected by one or more external workgroup keys, one or more shared keys, or any combination of workgroup keys and shared keys. Once the shares of secret information are encrypted, the first device may send the encrypted shares to the second device. Process flow 4400 then proceeds to step 4440.

At step 4440, the second device may attempt to recover the encrypted secret information. This recovery process may be dependent on how the shares of secret information were dispersed at step 4420 and encrypted at step 4430. For example, the secret information may have been dispersed into shares using an "M of N cryptosplit" and encrypted using a workgroup key at step 4420, and then those shares may be encrypted based on public keys of different certificate authorities at step 4430. The recovery process may decrypt the shares first using the public keys of the different certificate authorities, then decrypt the shares based on the workgroup key, and then use a restore function of the secure data parser to reconstruct the dispersed shares of secret information into the original secret information based on the "M of N cryptosplit".

If the recovery process is successful, the computed secret information may match the original secret information. This match may be mutually confirmed between the devices by each device independently computing a shared encryption key. For example, a first device may compute a shared encryption key based on the original secret information, while the second device computes a shared encryption key based on the recovered or restored secret information. In some embodiments, if the second device recovers the original secret information and subsequently computes a valid shared encryption key, process flow 4400 proceeds to step 4450, and messages are exchanged. In some embodiments, these messages may be securely exchanged based on the shared encryption keys computed by the first device and the second device. In some embodiments, if the second device does not recover the original secret information, process flow 4400 proceeds to step 4460, and messages are not exchanged. For example, messages may not be able to be exchanged because the shared encryption key of the first device does not match that of the second device.

FIG. 45 shows a simplified and illustrative process flow 4500 for establishing a secure proxy service between two devices, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. In process flow 4500, trust is distributed in a set of certificate authorities during initial negotiation of a connection between devices. In some embodiments, process flow 4500 may be executed as part of the key establishment phase of a secure exchange of information between two devices. This key establishment phase, including each of the steps in process flow 4500, may be part of one or more of the processes associated with the secure data parser, for example, similar to how the steps of process flow 4400 are associated with the secure data parser.

Process flow 4500 begins at step 4510. At step 4510, the devices exchange random numbers and certificates associated with public keys, each public key issued by a unique security authority. In some embodiments, these devices may be, for example, first device 4310 and second device 4320 as described with respect to process flow 4300 in FIG. 43. The exchanged random numbers may be generated by each device using the random number generator 3012 of secure data

parser 3026 as shown in FIG. 30. The exchanged certificates may be generated from the traversal of certificate authority hierarchy 4200 as shown in FIG. 42. For example, any suitable graph algorithm may traverse the certificate hierarchy 4200 to compute a list of the certificates of the root certificate authorities, or a list of the certificates of the minor certificate authorities within the tree of a single root certificate authority within certificate hierarchy 4200. In some embodiments, the exchanged certificates may be determined based on encryption parameters agreed upon by the first and second device. These parameters may be associated with the implementation of dispersal of the shares of secret information agreed upon by the first and second device. For example, if the dispersal techniques used at step 4520 disperses secret information into five shares, the exchanged certificates may include five certificates of unique certificate authorities from the first device, and five certificates of unique certificate authorities from the second device.

In some embodiments, a unique public key may be obtained for each unique certificate of the first device or the second device. In some embodiments, the encryption parameters may be set by the user of the secure proxy service, such as the user of a personal computer who wants to connect with the server of a financial institution using the secure proxy service. In some embodiments, the encryption parameters may be set by an administrator of the secure proxy service, such the administrator of the servers of a financial institution who want to offer secure proxy service enabled connections to their customers. In addition, in some embodiments, the exchanged lists of certificates may be based on an enrollment process 900 conducted with the user of one of the devices as described with respect to FIG. 9. Process flow 4500 then proceeds to step 4515.

At step 4515, the first device generates secret information. This secret information may be generated, for example, according to step 4410 described with respect to process flow 4400 of FIG. 44. Process flow 4500 then proceeds to step 4520. At step 4520, the first device disperses the secret information generated at step 4515 into shares using any suitable dispersal techniques. For example, the first device may perform an “M of N cryptosplit” of the secret information using the secure data parser of the present invention according to, for example, step 4420 described with respect to process flow 4400 of FIG. 44. In some embodiments, the shares of secret information resulting from the dispersal techniques may be restorable from at least a subset of the shares by recombining at least a quorum of the shares. In addition, in some embodiments, a keyed secret sharing routine may be applied to the shares of secret information using a keyed IDA. The key for the keyed IDA may be protected by one or more external workgroup keys, one or more shared keys, or any combination of shared and workgroup keys. Process flow 4500 then proceeds to step 4525.

At step 4525, the first device encrypts each share of secret information based on a public key issued by a different certificate authority. The public keys may be public keys obtained from the certificates sent to the second device from the first device at step 4510. In some embodiments, step 4525 may be included as part of step 4520. For example, the keyed secret sharing routine described with respect to step 4520 may be applied to the shares of secret information, where the keys for the keyed IDA are the public keys associated with the list of certificates sent to the second device from the first device. In another example, the public keys associated with the list of certificates may be used as split keys to encrypt the shares of secret information as described with respect to

options 3600 of FIG. 36. Process 4500 then proceeds to step 4530, or may optionally proceed to step 4527.

At optional step 4527, the first device may perform a key-wrap on the keys applied to the shares of secret information at step 4525. In some embodiments, the key wrap may be any suitable key encryption or key wrap algorithm. The key wrap may operate on the entire shares of dispersed secret information produced at step 4520. Alternatively, the dispersed shares may be additionally dispersed into several blocks, and the key wrap may operate on these blocks. Process flow 4500 then proceeds to step 4530.

At step 4530, the first device transmits the encrypted shares of secret information to the second device. This transmission may occur over any suitable communication channel, such as that described with respect to communication link 105 in FIG. 1. The first device in process flow 4500 then proceeds to step 4545, while the second device proceeds to step 4535.

At step 4535, the second device may attempt to decrypt the encrypted shares received from the first device. This decryption process may be based on how the shares of secret information were encrypted at step 4520 and step 4525. For example, at step 4520, shares of secret information may have been produced from dispersal techniques that produce shares of secret information that may be restorable from at least a subset of the shares by recombining at least a quorum of the shares. At step 4525, each dispersed share may have been encrypted using the public keys obtained by different certificate authorities corresponding to the second device’s certificates. At step 4527, the shares may have been additionally encrypted using a keywrap based on a workgroup key. Based on this encryption, at step 4535 the second device may first decrypt the encrypted shares of secret information based on the public keys issued by unique certificate authorities, then decrypt the shares of secret information based on the workgroup key of the keywrap applied at step 4527. It will be understood that beyond this particular example, any suitable number and type of decryption steps may be performed at step 4535. Process flow 4500 then proceeds to step 4540.

At step 4540, the second device may attempt to restore the original secret information based on the decrypted shares computed at step 4535. This restore process may be based on how the secret information generated at step 4515 was dispersed at step 4520. For example, at step 4520, shares of secret information may have been produced using dispersing functions of the secure data parser according to any of the techniques described with respect to FIG. 33, FIG. 35, and FIG. 36. Based on this dispersing, at step 4540 the second device may restore the original secret information from the split using restore functions of the secure data parser according to any of the techniques described with respect to FIG. 34. Process flow 4500 then proceeds to step 4545.

At step 4545, the first and second devices may independently compute a shared encryption key based on the exchanged random numbers, and computed or original secret information. For example, the first device may perform several digital signatures, one for each of its certificates in its list of certificates, using its own random number, the second device’s random number, and the secret information generated at step 4515. These digital signatures may then be used as input to a key generation function that computes the shared encryption key for the first device. The second device may perform similar digital signatures to compute its own shared encryption key, but use the decrypted secret information instead of the original secret information. Process 4500 then proceeds to step 4550.

In some embodiments, at step 4550, the first device and second device determine whether they agree on the shared

encryption keys independently computed at step 4545. In some embodiments, this agreement may be determined by the first and second device exchanging messages encoded with the shared encryption key. For example, the first device may send the second device a message encrypted with the shared encryption key. If the second device is able to decrypt the encrypted message and respond, for example, with an appropriate acknowledgment, the first device may determine that it may securely exchange messages with the second device. Otherwise, the first device may determine that it may not securely exchange messages with the second device, and no further messages are exchanged. It will be understood that a similar determination may occur at the second device. In some embodiments, the first and second devices may determine that their independently computed shared encryption keys match by exchanging the encryption keys without any message. If the first and second devices do not agree on the shared encryption key, process 4500 proceeds to step 4555. If the first and second devices agree on the shared encryption key, process 4500 proceeds to step 4560. At step 4555, the first and second devices do not exchange messages. At step 4560, the first and second devices exchange messages. After each of steps 4555 and 4560, process flow 4500 may end.

FIG. 46 shows a simplified and illustrative process flow 4600 for establishing a secure proxy service between client 4610 and server 4620, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. In process flow 4600, trust is distributed in a set of certificate authorities during initial negotiation of a connection between client 4610 and server 4620. In some embodiments, process flow 4600 may be executed as part of the key establishment phase of a secure exchange between client 4610 and server 4620. This key establishment phase, including any of the steps in process flow 4600, may be part of one or more of the processes associated with the secure data parser according to, for example, how the steps of process flow 4400 are associated with the secure data parser. In addition, process flow 4600 may be an example of process flow 4500 as discussed with respect to FIG. 45.

Process flow 4600 begins at step 4612. At step 4612, client 4610 sends server 4620 a generated random number R_C and a list of its certificates Cert-CA1_C , Cert-CA2_C , and Cert-CA3_C . Client 4610 and server 4620 may be any suitable client and server devices as described with respect to user system 105 and vendor system 120 of FIG. 1, respectively. R_C may be generated by client 4610, according to, for example, how the random number generated by the first device at step 4510 of process flow 4500 in FIG. 45. Each of these certificates may be associated with a public key issued by a different security authority, similar to the lists of certificates discussed with respect to step 4510 of process flow 4500 in FIG. 45. Process flow 4600 then proceeds to step 4622.

At step 4622, server 4620 sends client 4610 its own generated random number R_S and a list of its certificates Cert-CA1_S , Cert-CA2_S , and Cert-CA3_S . R_S may be generated by server 4620 according to the random number generated by the second device at step 4510 of process flow 4500 in FIG. 45. Each of these certificates may be associated with a public key issued by a unique security authority, similar to the public keys of the unique certificate authorities discussed with respect to step 4510 of process flow 4500 in FIG. 45. Process flow 4600 then proceeds to step 4614.

At step 4614, client 4610 generates secret information. This secret information may be generated according to, for example, step 4515 of process flow 4500 in FIG. 45. Also at step 4614, client 4610 disperses the secret information S_C into

shares $S1_C$, $S2_C$, and $S3_C$. This dispersing may be performed according to, for example, the dispersing of secret information discussed with respect to step 4520 of process flow 4500 in FIG. 45. Also at step 4614, client 4610 encrypts the shares of secret information using a different one of server's public keys. For example, if "Enc" represents the encryption function executed by the secure data parser, and Pub1_S , Pub2_S , and Pub3_S represent the public keys corresponding to the server's certificates Cert-CA1_S , Cert-CA2_S , and Cert-CA3_S , respectively, the client may encode $S1_C$ using Pub1_S by executing $\text{Enc}(\text{Pub1}_S, S1_C)$, may encode $S2_C$ using Pub2_S by executing $\text{Enc}(\text{Pub2}_S, S2_C)$, and may encode $S3_C$ using Pub3_S by executing $\text{Enc}(\text{Pub3}_S, S3_C)$. The encryption function may be chosen out of any combination of the methods of encryption described with respect to steps 4525 and 4527 of process flow 4500 in FIG. 45. Once the secret information is generated, dispersed, and encrypted, the encrypted shares are transmitted to server 4620. Process flow 4600 then proceeds to steps 4616 and 4624.

At step 4624, the shares of secret information $S1_C$, $S2_C$, and $S3_C$ may be decrypted and restored into the original secret information by server 4620 using any suitable decryption and restoring techniques described with respect to steps 4525 and 4540 of process flow 4500 in FIG. 45. Server 4620 may then use the restored secret information to generate a shared encryption key K using a key generation function G . Key generation function G may take random numbers R_C and R_S as input along with the restored secret information. At step 4616, client 4610 may similarly generate its own shared encryption key K using a key generation function G . However, the key generation function executed by client 4610 may use the original secret information generated by client 4610 rather than the restored secret information generated by server 4620. Process 4600 then proceeds to step 4630.

At step 4630, messages are exchanged between client 4610 and server 4620 using their respective shared encryption keys K . In some embodiments, client 4610 and server 4620 may determine whether their shared encryption keys match similar to the process described with respect to step 4550 of process flow 4500 in FIG. 45. If it is determined that the respective shared encryption keys of client 4610 and server 4620 do not match, messages may not be exchanged or may cease to be exchanged between client 4610 and server 4620. Otherwise, the exchange of messages may continue similar to normal TLS or SSL communication after the key establishment phase.

In some embodiments, a secure proxy service may be resident on a client application running on client 4610. The client application may maintain a list of secure proxy server enabled servers, such as server 4620, based on the IP address or URL and port number of the servers. In some embodiments, the client application may be associated with an address that is addressable by the servers. When a connection is requested by the client for a secure proxy service enabled server, the client application may establish a connection with the specified server proxy service, utilizing the approaches described in process flows 4400, 4500, and 4600. In addition, a secure proxy service may be resident on a server application running on service 4620. The server application may accept connections from the client application, and forwards the data it receives to the proper secure proxy configured port based on port forwarding rules. These port forwarding rules may be predetermined or mutually agreed upon by the client application and the server application.

The key establishment phase described by process flows 4400, 4500, and 4600 offer the assurance that if some, but less than a quorum, of the certificate authorities have been com-

promised, the connection can still be securely established between two devices. That is, even if compromised certificate authorities have access to the information exchanged between the devices they would not have enough information to disrupt the confidentiality or integrity of the communication. For example, if there were three shares of secret information each encrypted with a public key of a different certificate authority as shown in process flow **4600**, one of the certificate authorities could be compromised and the connection could be securely established between two devices. This security is ensured because even if the compromised certificate authority had access to the messages being passed between two devices, the attacker associated with the compromised certificate authority would not have knowledge of the public and private key pair of the other two certificate authorities, and thus would at most be able to recover one of the shares of secret information. Further, because the shares of secret information were dispersed such that they could be restored with at least a subset of the shares by recombining at least a quorum of the shares, the attacker behind the compromised certificate authority would not be able to construct the original secret information using just one recovered share. Accordingly, the attacker behind the compromised certificate authority would not be able to recover the secret information, and would not be able to compute the shared encryption key used to encrypt messages between the first device and the second device.

Process flows **4400**, **4500**, and **4600** are described in various embodiments as occurring between two devices that wish to establish a secure means of communication between them. However, in some embodiments process flow **4400** may occur between more than two devices. For example, process flow **4400** may be used to establish a secure means of communication between a personal computer, and a set of web servers. Each web server in the set may use a different set of unique certificate authorities in the key establishment phase of communication with the first device.

FIG. **47** and FIGS. **48A** and **48B** show simplified and illustrative process flows **4700**, **4800**, and **4850** for establishing a secure proxy service between devices, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. In process flows **4700**, **4800**, and **4850**, trust is distributed among a set of certificate authorities in the structure of the communication channel used to exchange messages between the devices. In some embodiments, process flow **4700** may be executed after the key establishment phase of a secure exchange of information between devices, but before the devices exchange messages. Process flows **4700**, **4800**, and **4850** may be part of one or more of the processes associated with secure data parser **3706** as illustrated in FIG. **37**, or may be a standalone process. For example, steps **4710**, **4720**, **4730**, **4740**, **4750**, and **4760** may be part of the post-encryption process **3720** associated with secure data parser **3706** as illustrated in FIG. **37**, or may be a standalone process.

Process flow **4700** begins at step **4710**. At step **4710**, a communication channel is established between devices. This communications channel may be established using any suitable trust engine **110** described with respect to FIG. **1** through FIG. **14**. In some embodiments, this communication channel may be secured using any suitable encryption technology to secure data in motion in any suitable communications. For example, the communication channel may be established using conventional SSL, $\frac{1}{2}$ SSL, FULL SLL, TLS, Full TLS, RS1, OTP, RC4™, Triple DES, AES, IPsec, public key encryption, symmetric key encryption, split key encryption, multi-factor encryption, or any suitable combination of

encryption technologies. In some embodiments, this communication channel may not be secure. For example, the established communication channel may carry data through clear text. In some embodiments, these encryption technologies may use keys issued from a certificate authority. This certificate authority may be referred to as an “outer level certificate authority” because it may secure the first communication channel independently of any certificate authorities used to secure the secure communication tunnels described with respect to step **4730** below.

In addition, the communication channel may carry data associated with email, streaming data broadcasts, and WiFi communications. In some embodiments, the established communication channel may utilize any number of server or client-side technologies, such as CGI scripts, ASPs, or any suitable web server technologies. In some embodiments, the communication channel may be established across several physical transport mediums or physical paths. For example, the communication channel may be established over one or more of the Internet, an intranet, a LAN, WiFi, Bluetooth, WiMax, or any suitable hard-wired or wireless communication means, or any combination thereof. Each physical transport medium may have a different network topology between the devices that exchange messages on the particular physical medium. Process **4700** then proceeds to step **4720**.

At step **4720**, any number of secure communication tunnels are established within the first communication channel based on distributed trust among a set of certificate authorities. These certificate authorities may be referred to as “inner level certificate authorities” because they may protect the communications within the secure communication tunnels independently of any outer level certificate authority. In some embodiments, these communication tunnels may be established using any suitable key establishment phase of any of the encryption technologies described with respect to step **4710**. In some embodiments, the secure communication tunnels are established using an encryption technology that is different from that utilized by the first communication channel. For example, the communication channel may be established using AES, while the secure communication tunnels are established using full TLS. In this example, each of the secure communication tunnels may be established using a key establishment process similar to that described with respect to process flow **4300** of FIG. **43**. In some embodiments, the secure communication tunnels are established using the same encryption technology as the first communication channel. For example, the communication channel and each of the secure communication tunnels may be established using full TLS.

In some embodiments, the secure communications tunnels may be established using the same encryption technology, for example, each communication channel may be established using full TLS. In other embodiments, the secure communication tunnels may be established using different encryption technologies, for example some of the encryption tunnels may be established using full TLS, while other tunnels are established using AES. In some embodiments, the secure communication tunnels may be established across several physical media or physical paths. For example, the secure communication tunnels may be established over one or more of the Internet, an intranet, a LAN, WiFi, Bluetooth, WiMax, or any suitable hard-wired or wireless communication means, or any combination thereof. Each physical transport medium may have a different network topology between the devices that exchange messages on the particular physical medium.

Regardless of which encryption technologies are used to establish the secure communication tunnels, the tunnels are

established at step 4720 based on distributed trust among certificate authorities. In some embodiments, this distributed trust may be achieved by establishing each secure communication tunnel based on a unique certificate authority. In some embodiments, each secure communication tunnel may be established using a certificate issued one of the unique certificate authorities. In such embodiments, symmetric encryption key material may be communicated during the establishment of each channel using the certificate issued by the unique certificate authority associated with that channel. In such embodiments, the symmetric key encryption material may be, for example, the symmetric encryption key material discussed with respect to the use of certificate authorities in full TLS above. Each certificate authority may be unique in that the certificates issued by each have different public and private key pairs. In some embodiments, the unique certificate authorities may be root certificate authorities. In other embodiments, the unique certificate authorities may be minor certificate authorities within the tree of a single root certificate authority, as discussed with respect to certificate authority hierarchy 4200 in FIG. 42. The unique public and private key pairs of the different certificate authorities may be used during key establishment of each secure communication tunnel. For example, if the secure communication tunnels are based on TLS, each of the tunnels may be established as described with respect to process flow 4300 of FIG. 43 using the certificate of one of the unique certificate authorities. Process 4700 then proceeds to step 4730.

At step 4730, data packets are prepared for transmission over the secure communication tunnels based on the set of certificate authorities and multi-factored secret sharing. In some embodiments, this preparation may include encrypting the data packets using a key developed during the establishment of a different one of the communication tunnels. In some embodiments, this key may be a symmetric key generated using symmetric encryption key material that was communicated during the establishment of each channel using a certificate of a unique certificate authority associated with that channel. In addition, this preparation may include dispersing each data packet into shares based on multi-factored secret sharing, and then encrypting the resulting shares using the certificates of the unique certificate authorities used to establish the secure communication tunnels at step 4720. This dispersing process may be achieved using any suitable data splitting or cryptosplit as discussed with respect to data splitting module 520 or 610 of FIG. 5 and FIG. 6, and elaborated with respect to FIG. 8 and FIGS. 20 through 24.

In addition, in some embodiments, the data packets may be dispersed into shares substantially through any number of uses of the secure data parser outlined with respect to FIG. 33, FIG. 35, and FIG. 36. For example, the secure data parser may receive unencrypted data packets. If the data packets are going to be split with an algorithm that requires a key, a split encryption key is generated. In some embodiments, the data packets may be split into shares at step according to any of the techniques described with respect to FIG. 33, FIG. 35, and FIG. 36. In some embodiments, a fault tolerant scheme may be used to encrypt the split encryptions key and allow for regeneration of the data packets from less than the total number of shares. For example, the data packets may be dispersed into shares using any suitable data dispersion techniques such that the shares are restorable from at least a subset of the shares by recombining at least a quorum of the shares. In addition, information may be added to the shares of data packets that are used to reconstruct the data packets. Further, once the shares are created, authentication/integrity information may be embedded into the shares of data packets. Each

share may be post-encrypted using public keys of unique certificate authorities used to establish the secure communication tunnels at step 4720.

In addition, in some embodiments, the dispersing of the data packets may occur in two phases—a header generation phase and a data partitioning phase. The phases may be, for example, the simplified header generation process 4000 as shown in FIG. 40A and simplified data partitioning process 4010 as shown in FIG. 40B. One or both of these processes may be performed on the data packets. In some embodiments, the data packets may be pre-encrypted based on the encryption technology used to establish the first communication channel. The pre-encrypted data packets may then be run through processes 4000 and/or 4010 as described below.

As shown at step 4002 of FIG. 40A, split keys may be generated. The unencrypted or pre-encrypted data packets may then be optionally encrypted (e.g., using the workgroup key described with respect to FIG. 39) at step 4004. At step 4006, a “M of N cryptosplit” may be used to split the data packets into shares of secret information using the split key. Each share of the data packet may then be embedded into a share header. Finally, a share integrity block and a post-authentication tag (e.g., MAC) may be appended to the header block of each share. Each header block may be designed to fit within a single encrypted data packet. Each header block may be post-encrypted using a key developed during the establishment of a different one of the communication tunnels at step 4720.

In some embodiments, after the headers including the shares of secret information are generated, the secure data parser may enter a data partitioning phase. This data partitioning phase may be, for example, the simplified data splitting process 4010 as shown in FIG. 40B. For example, each incoming unencrypted or pre-encrypted data packet may be encrypted using one or more keys, such as a shared key or a workgroup key, at step 4012. In some embodiments, the data that is encrypted may include the headers that contain the shares of data packets computed during simplified header generation process 4000. At step 4014, share integrity information (e.g., a hash H) may be computed on the resulting ciphertext from step 4012. For example, a SHA-256 hash may be computed on the data that is encrypted with one or more keys at step 4012. At step 4106, the data packet may then be partitioned into two or more data shares using one of the data splitting algorithms described above in accordance with the present invention. In some embodiments, the data packet or data block may be split so that each data share contains a substantially random distribution of the encrypted data packet or data block. The integrity information (e.g., hash H) may then be appended to each data share. An optional post-authentication tag (e.g., MAC) may also be computed and appended to each data share in some embodiments. In addition, each data share may include metadata as described with respect to FIG. 40B. The metadata may include information regarding the data packets and workgroup keys used to generate the shares of data packets. Each data packet share may be post-encrypted using public keys of unique certificate authorities used to establish the secure communication tunnels at step 4720.

In some embodiments, the shares of data packets may be associated with shares of an encryption key or data split key similar to how key and data components are stored within shares as shown in illustrative block diagrams 3800 and 3900 in FIGS. 38 and 39. For example, shares of the data packets may be stored similarly to how portions of encryption key 3804 are split and stored within shares 3810. When a workgroup key is utilized, the shares of data may be encrypted with

the workgroup key prior to be stored within the shares as shown in illustrative block diagram 3900.

Regardless of how the data packets are dispersed into shares, each share may be post-encrypted using a key developed during the establishment of a different one of the communication tunnels at step 4720. For example, in some embodiments there may be three secure communication tunnels established at step 4720, each with a unique certificate authority. Each share produced at step 4730 may then be encrypted using the key developed during the establishment of a different one of the communication tunnels. In some embodiments, these keys may be the symmetric keys generated using symmetric encryption key material that was communicated during the establishment of each channel using the three unique certificate authorities. Process 4700 then proceeds to step 4740.

At step 4740, the prepared data packets are transmitted to their destination. This transmission may occur over any suitable communication channel, such as that described with respect to communication link 105 in FIG. 1. In some embodiments, the destination for the packets may be one or more of secure data proxy enabled servers. A client application running on a client device may maintain a list of secure proxy server enabled servers based on the IP address or URL and port number of the servers. In some embodiments, the client application may be associated with an address that is addressable by the servers. When a connection is requested by the client for a secure proxy service enabled server, the client application establishes a connection with the specified server proxy service, utilizing the approaches described in steps 4710 and 4720. Once the prepared data is transmitted to its destination, process 4700 then proceeds to step 4750.

At step 4750, the transmitted data packets are received. In some embodiments, the transmitted data packets may be received by a secure proxy service that is resident on a server application running on a secure data proxy enabled server. The server application may accept connections from the client application, and forwards the data it receives to the proper secure proxy configured port based on port forwarding rules. These port forwarding rules may be predetermined or mutually agreed upon by the client application and the server application. Process 4700 then proceeds to step 4760.

At step 4760, the data packets are restored based on the set of certificate authorities and the multi-factored secret sharing. In some embodiments, this restoration may be a mirrored process of the preparation process used to parse and encrypt the data packets at step 4730. For example, at step 4730, the data packets may have been encrypted using keys associated with the establishment of the secure communication tunnels at step 4720. In addition, shares of unencrypted or pre-encrypted data packets may have been produced using dispersing techniques of the secure data parser according to any of the techniques described with respect to FIG. 33, FIG. 35, and FIG. 36. In some embodiments, the shares of data packets may be encrypted based on the encryption technology used to establish the first communication channel.

Accordingly, the shares of data packets may first be decrypted based on the keys associated with the establishment of the secure communication tunnels at step 4720. The decrypted shares may then be restored using restore functions of the secure data parser according to any of the techniques described with respect to FIG. 34. In some embodiments, the restored shares may be decrypted based on the encryption technology used to establish the first communication channel. In some embodiments, the decrypted shares of data may be embedded in share headers. In such embodiments, the shares of data may be extracted from the decrypted share headers,

and restored using the restore functions of the secure data parser. Process flow 4700 then ends. In some embodiments, steps 4730, 4740, 4750, and 4760 may be repeated as necessary for the transmission of data over the secure communication tunnels.

Describing embodiments of the secure proxy service with respect to FIG. 48A, process flow 4800 may be executed on a first device, such as on a client-side application running on a personal computer that requests to communicate using the secure proxy service with a second device, such as a web server. Process flow 4800 begins at step 4810. At step 4810, a first secure communication channel may be established. This secure communication channel may be established using a key establishment process with the keys of any suitable encryption technologies as described with respect to step 4710 in process flow 4700 of FIG. 47. In some embodiments, these keys may be issued from a certificate authority referred to as an "outer level certificate authority". In addition, the communication channel may carry data associated with any suitable applications as discussed with respect to step 4710 in process flow 4700 of FIG. 47. Also, the communication channel may be established across several physical transport mediums as described with respect to step 4710 of process flow 4700 of FIG. 47. Process 4800 then proceeds to step 4815.

At step 4815, any number of secure communication tunnels are established within the first communication channel (e.g., one, two, three, five, ten, fifty, one hundred, or more than one hundred secure communication tunnels). Each secure communication channel may be established using a certificate obtained from a unique certificate authority and each tunnel may be associated with the respective unique certificate authority. In some embodiments, symmetric encryption key material may be communicated during the establishment of each channel using the certificate issued by the unique certificate authority associated with that channel. In such embodiments, the symmetric key encryption material may be, for example, the symmetric encryption key material discussed with respect to the use of certificate authorities in full TLS above. Similar to process flow 4700 of FIG. 47, the unique certificate authorities used to establish the secure communication tunnels may be referred to as "inner level certificate authorities". Each of the secure communication tunnels may be established using a key establishment process with any suitable encryption technologies over one or more physical transport media as described with respect to step 4720 of process flow 4700 of FIG. 47. Also similar to step 4720 of process flow 4700 of FIG. 47, each certificate authority may be unique in that the certificates issued by each have different public and private key pairs. In some embodiments, each secure communication tunnel may be associated with a respective unique certificate authority in that all data sent over that tunnel is encrypted based on keys developed during the establishment of the communication tunnels. In some embodiments, this association may be tracked in any suitable data structure by the secure data parser at a client application, a server application, or both. Process 4800 then proceeds to step 4820.

At step 4820, incoming data packets may be cryptographically split into any number of shares using multi-factored secret sharing. In some embodiments, the incoming data packets may be split into the same number of shares as the number of secure communication tunnels established at step 4815. The cryptographic split of the incoming data packets may be achieved according to, for example, any suitable

dispersing techniques discussed with respect to step 4730 of process flow 4700 of FIG. 47. Process 4800 then proceeds to step 4825.

At step 4825, each of the shares is encrypted using a key developed during the establishment of a different one of the secure communication tunnels. In some embodiments, the key may be a symmetric encryption key generated using symmetric encryption key material that was communicated during the establishment of each channel using a certificate of a unique certificate authority associated with that channel. In some embodiments, this symmetric encryption key material may be, for example, the symmetric encryption key material discussed with respect to the use of certificate authorities in full TLS above. It may be understood, however, that the keys developed during the establishment of the secure communication tunnels may be any suitable encryption key, secret information, or any other information other than symmetric encryption keys. For example, the keys developed during the establishment of the secure communication tunnels may be asymmetric encryption keys. In some embodiments, each of the shares that are produced at step 4820 is preprocessed and tagged with one or more bits that identify which of the keys associated with the establishment of the communication tunnels should be used to encrypt each of the shares. Process 4800 then proceeds to step 4830.

At step 4830, each of the encrypted shares is transmitted on the tunnel associated with the unique certificate authority under which that tunnel was established. For example, if there were three secure communication tunnels established each based on a different one of three unique certificate authorities, each incoming unencrypted or pre-encrypted data packet would be cryptographically split into three shares using multi-factored secret sharing and encrypted using a different one of three keys developed during the establishment of the three secure communication tunnels using a different one of three unique certificate authorities. Accordingly, each one of the three encrypted shares would be transmitted on the tunnel associated with the unique certificate authority under which that tunnel was established. In some embodiments, this transmission may be based on the data structure that maintains the associations between the certificate authorities and the tunnels.

In some embodiments, the association between a certificate authority and the secure communication tunnels may remain constant throughout the duration of a data transmission. In other embodiments, the associations between the certificate authorities and the secure communication tunnels may be dynamic. In such embodiments, the associations may be shuffled at any suitable point in time, such as after the transmission of an entire data packet. For example, a first data packet may be processed by process flow 4800 wherein the data packet is cryptosplit into three shares at step 4820. The first share of the first data packet may be encrypted using a first key developed during the establishment of a first secure communication tunnel using a first certificate authority and transmitted over the first communication tunnel. The second share of the first data packet may be encrypted using a second key developed during the establishment of a second secure communication tunnel using a second certificate authority and transmitted over the second communication tunnel. Finally, the third share of the first data packet may be encrypted using a third key developed during the establishment of a third secure communication tunnel using a third certificate authority and transmitted over the third communication tunnel.

In some embodiments, after the three shares of the first data packet are transferred, the associations between the certificate authorities and the communication channels may be shuffled

such that the first share may be encrypted using the third key and transmitted over the third tunnel, the second share may be encrypted using the first key and transmitted over the first tunnel, and the third certificate authority may be encrypted using the second key and transmitted over the second tunnel. In such embodiments, these associations may be stored at any suitable depository that is accessible to the devices in communication, such as depository 210 of FIG. 2.

Describing embodiments of the secure proxy service with respect to FIG. 48B, process flow 4850 may be executed on a client-side secure proxy server application running on a second device, such as a web server, that is exchanging information with a first device, such as a personal computer running a client-side secure proxy service application. Process flow 4850 begins at step 4835. At step 4835, encrypted shares of data are each received on a respective communication tunnel. The server application may accept connections from the client application, and forward the data it receives to the proper secure proxy configured port based on port forwarding rules. These shares of data may be the same shares that were cryptographically split, encrypted, and transmitted at steps 4820, 4825, and 4830, of process flow 4800 of FIG. 48A, respectively. Process flow 4850 then proceeds to step 4840.

At step 4840, each of the shares are decrypted based on the key associated with the establishment of the respective secure communication tunnel that the share was received on. In some embodiments, this process may be mirrored of that described with respect to step 4825 of process flow 4800 of FIG. 48. Process flow 4850 then proceeds to step 4845.

At step 4845, the decrypted data packet shares are restored into the original data packets. In some embodiments, this restoration may be a mirrored process of the dispersing techniques used at step 4820 of process flow 4800 in FIG. 48. In some embodiments, the restored shares may be decrypted based on the data dispersion and/or encryption technologies used to establish the first communication channel. Process flow 4800 then ends. In some embodiments, the steps of process flows 4800 and 4850 may be repeated as necessary for the transmission of data over the secure communication tunnels.

The communication protocols described with respect to process flows 4700, 4800, and 4850 offer the assurance that if certain outer or inner level certificate authorities have been compromised, data will be securely exchanged between devices. That is, even if compromised certificate authorities have access to the information exchanged based on the keys associated with that certificate authority, the attacker associated with the compromised certificate authority would not have enough information to disrupt the confidentiality or integrity of the communication. For example, if the outer level certificate authority was compromised but the inner level certificate authorities retained their integrity, the attacker would be able to observe the streams of data within each of the secure communication tunnels. However, the attacker would have no knowledge of the encryption used within each of the secure communication tunnels, including knowledge of the certificate of each of the unique certificate authorities used to secure data over each secure communication tunnel.

In another example, if one or more of the inner level certificate authorities was compromised but the outer level certificate authority remained intact, the attacker may be able to recover cryptographically split portions of the data packets, but may not be able to decrypt the cryptographically split portions because it would have no knowledge of the encryption used by the outer level certificate authority. Further, if the data packets sent through the secure communication tunnels

are cryptographically split such that they are restorable from at least a subset of the shares by recombining at least a quorum of the shares, the user of the secure data parser may have the additional protection that if some, but less than a quorum, of the certificates associated with the secure communication tunnels have been compromised, the attacker would not be able to restore the cryptosplit data packets.

FIG. 48C is a simplified block diagram of a secure proxy service 4870 that distributes trust among a set of certificate authorities in the structure of communication channels, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. Secure proxy service 4870 may reside on any suitable trust engine 110 or module within trust engine 110 as described with respect to FIGS. 1-8. Secure proxy service is illustrated as including a first communication channel 4880 and sub-channels (i.e., secure communication tunnels) 4872, 4874, and 4876 that are established using full TLS. However, it will be understood that these communication channels may be established and used with any suitable encryption technologies, or without encryption, as discussed with respect to steps 4710 and 4720 of process flow 4700 in FIG. 47, or as discussed with respect to steps 4810 and 4815 of process flow 4800 of FIG. 48A. In addition, although secure proxy service 4870 is illustrated as using three sub-channels, any suitable number of sub-channels may be used to securely transfer information over the secure proxy service 4870.

Secure proxy service 4870 may include received data packets 4877. In some embodiments, data packets 4877 may be unencrypted packets of data to be processed by secure proxy service 4870. Data packets 4877 may be received from any suitable source, such as depository 210 described with respect to trust engine 210 of FIG. 2. In other embodiments, data packets 4877 may be pre-encrypted according to the encryption technology used to establish communication channel 4880. For example, as shown in secure proxy service 4870, communication channel 4880 is established using full TLS. This full TLS communication channel may be established according to process flow 4300 of FIG. 43.

Secure proxy service 4870 may also include cryptographic packet split module 4878. In some embodiments, cryptographic packet split module 4878 may include any circuitry and/or instructions for executing and/or computing any of the encryption and data dispersing techniques discussed with respect to step 4730 of process flow 4700 of FIG. 47, or discussed with respect to step 4820 and 4825 of process flow 4800 of FIG. 48A. In some embodiments, cryptographic packet split module 4878 may reside on a client device or a client-side application that requests to communicate with a server. In other embodiments, cryptographic packet split module 4878 may reside or run on any device that is suitable to run secure data parser 3706 of illustrative overview process 3700 of FIG. 37.

Secure proxy service 4870 may also include communication channel 4880. Communication channel 4880 may be established over one or more physical mediums using any suitable encryption technologies, or no encryption, as described with respect to the first communication channel at step 4710 of process flow 4700 of FIG. 47, or as described with respect to first secure communications channel at step 4810 of process flow 4800 of FIG. 48A. Sub-channels 4872, 4874, and 4876 may be established based on communication channel 4880. These sub-channels may be established over one or more physical mediums based on a certificate of a unique certificate authority according to, for example, the secure communication tunnels are described with respect to

step 4720 of process flow 4700 of FIG. 47, or are described with respect to step 4815 of process flow 4800 of FIG. 48A. In this manner, each sub-channel may be associated with a unique certificate authority.

For example, as shown in secure proxy service 4870, TLS sub-channel 4872 is associated with certificate authority CA1, TLS sub-channel 4874 is associated with certificate authority CA2, and TLS sub-channel 4876 is associated with certificate authority CA3. In some embodiments, the associations between sub-channels and their respective certificate authorities may remain constant. In other embodiments, the associations between sub-channels and their respective certificate authorities may change as described with respect to step 4830 of process flow 4800 of FIG. 48A. Communication channel 4880 and sub-channels 4872, 4873, and 4876 may be established on any suitable communication link, such as communication link 125 described with respect to cryptographic system 100 of FIG. 1.

In some embodiments, cryptographic packet split module 4878 may transmit dispersed shares of data packets 4877 over sub-channels 4872, 4874, and 4876. This transmission may occur according to, for example, step 4740 of process flow 4700 of FIG. 47, or according to, for example, step 4830 of process flow 4800 of FIG. 48A. In some embodiments, cryptographic packet split module 4878 may transmit one of the split shares over each of the sub-channels 4872, 4874, and 4876. In other embodiments, cryptographic packet split module 4878 may transmit more than one split share over one or more of the sub-channels 4872, 4874, and 4876. Such embodiments may be useful when one of the sub-channels 4872, 4874, and 4876 is unusable due to a failure in the physical medium supporting one of the sub-channels.

In some embodiments, sub-channels 4872, 4874, and 4876 may include data encryption modules 4871, 4873, and 4875, respectively. Data encryption modules 4871, 4873, and 4875 may each be associated with a unique certificate authority associated with one of the sub-channels. In some embodiments, data encryption modules will apply data encryption to each share of a data packet that passes over the sub-channel. For example, as shown with respect to proxy service 4870, encryption module 4871 is associated with certificate authority CA1 that is associated with TLS sub-channel 4872, and encrypts each share of a data packet that passes over sub-channel 4872 using full TLS based on a key developed during establishment of channel 4872. The establishment of channel 4872 may have used a certificate obtained from certificate authority CA1. Encryption module 4873 is associated with certificate authority CA2 that is associated with TLS sub-channel 4874, and encrypts each share of a data packet that passes over sub-channel 4874 using full TLS based on a key developed during establishment of channel 4874. This establishment of channel 4874 may have used a certificate obtained from certificate authority CA2. Finally, encryption module 4875 is associated with certificate authority CA3 that is associated with TLS sub-channel 4876, and encrypts each share of a data packet that passes over sub-channel 4876 using full TLS based on a key developed during establishment of channel 4876. This establishment of channel 4876 may have used a certificate obtained from certificate authority CA3. In some embodiments, encryption modules 4872 may reside or run on any device that is suitable to run secure data parser 3706 of illustrative process 3700 of FIG. 37.

Secure proxy service 4870 may also include packet restore module 4879. Packet restore module may receive shares of data packets from sub-channels 4872, 4874, and 4876 as described with respect to step 4750 of process flow 4700 of FIG. 47, or as described with respect to step 4835 of process

flow 4850 of FIG. 48B. In some embodiments, packet restore module 4879 may include any circuitry and/or instructions for executing and/or computing any of the decryption techniques or packet restoration techniques to produce restored data packets 4882 as described with respect to step 4760 of process flow 4700 of FIG. 47 or steps 4840 and 4845 as described with respect to process flow 4850 of FIG. 48B. In some embodiments, restore packet module 4879 may reside on a server device or a service-side application that receives requests from a client device or client-side application. In other embodiments, packet restore module 4879 may reside or run on any device that is suitable to run secure data parser 3700 of illustrative overview process 3700 of FIG. 37.

FIG. 49 is a process flow diagrams of illustrated steps and features for a secure proxy service 4900 between client 4910 and server 4920 that distributes trust among a set of certificate authorities in the structure of communication channels, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. For example, trust may be distributed in first certificate authority CA1, second certificate authority CA2, and third certificate authority CA3 as part of key establishment phase 4570, as will be described below. In some embodiments, secure proxy service 4900 may be executed during and after the key establishment phase of a secure exchange of information between client 4910 and server 4920, but before the client 4910 and server 4920 exchange messages. Secure proxy service 4900 may be part of one or more of the processes associated with the secure data parser similar to how the steps of process flows 4700, 4800, and 4850 are associated with the secure data parser. In addition, secure proxy service 4900 may be an example of process flows 4700, 4800, and 4850 of FIGS. 47, 48A, and 48B, or may be an example of the operation of secure proxy service 4870.

Secure proxy service 4900 begins at step 4930. At step 4930, client 4910 establishes a first communication channel (not shown) and secure communication tunnels as described with respect to first communication channel and secure communication tunnels at steps 4710 and 4720 of process flow 4700 of FIG. 47, and first secure communication channel and secure communication tunnels at steps 4810 and 4820 of process flow 4800 of FIG. 48A. Secure proxy service 4900 then proceeds to step 4912.

At step 4912, client 4910 may generate secure data parser headers H_1 , H_2 , and H_3 , and transmit them to server 4920. Headers H_1 , H_2 , and H_3 may contain information related to the data dispersion techniques agreed upon by client 4910 and server 4920. For example, in some embodiments, client 4910 and server 4920 may agree upon the use of an "M of N cryptosplit" of each exchanged message. Headers H_1 , H_2 , and H_3 may be encrypted based on keys associated with the establishment of the secure communication tunnels at step 4930. In some embodiments, this encryption and header generation may be included as part of any of the encryption and data dispersing techniques discussed with respect to step 4730 of process flow 4700 of FIG. 47, or discussed with respect to step 4820 and 4825 of process flow 4800 of FIG. 48A. In addition, in some embodiments, this header generation process may be executed by cryptographic packet split module 4878 as described with respect to secure proxy service 4870 of FIG. 48C. Client 4910 may then transmit Headers H_1 , H_2 , and H_3 to server 4920. This transmission may occur according to, for example, step 4740 of process flow 4700 of FIG. 47, or according to, for example, step 4830 of process flow 4800 of FIG. 48A. In addition, a cryptographic packet split module 4878 may transmit Headers H_1 , H_2 , and H_3 of data

packets 4877 over sub-channels 4872, 4874, and 4876 as described with respect to secure proxy service 4870 in FIG. 48C. Secure proxy service 4900 then proceeds to step 4922.

At step 4922, server 4920 may receive Headers H_1 , H_2 , and H_3 as described with respect to step 4750 of process flow 4700 of FIG. 47, or as described with respect to step 4835 of process flow 4850 of FIG. 48B. Client 4910 may then disperse data D into shares D_1 , D_2 , and D_3 . Shares D_1 , D_2 , and D_3 may be encrypted based on keys associated with the establishment of the secure communication tunnels at step 4930. In some embodiments, this dispersing process may be included as part of any of the encryption and data dispersing techniques discussed with respect to step 4730 of process flow 4700 of FIG. 47, or discussed with respect to step 4820 and 4825 of process flow 4800 of FIG. 48A. Secure proxy service 4900 transmits the shares of encrypted and parsed data to client 4910. This transmission may occur according to, for example, step 4740 of process flow 4700 of FIG. 47, or according to, for example, step 4830 of process flow 4800 of FIG. 48A. Secure proxy service 4900 may proceed to step 4914.

At step 4914, client 4910 may receive shares D_1 , D_2 , and D_3 as described with respect to step 4750 of process flow 4700 of FIG. 47, or as described with respect to step 4835 of process flow 4850 of FIG. 48B. In some embodiments, a packet restore module may receive shares of data packets from sub-channels as described with respect to packet restore module 4879 of secure proxy service 4870 in FIG. 48C. Client 4910 may then decrypt and restore shares D_1 , D_2 , and D_3 . Shares may be decrypted and restored according to any suitable decrypt and restore techniques as described with respect to step 4760 of process flow 4700 of FIG. 47 or steps 4840 and 4845 as described with respect to process flow 4850 of FIG. 48B. In some embodiments, shares D_1 , D_2 , and D_3 may be decrypted and restored by a packet restore module such as packet restore module 4879 of secure proxy service 4870 in FIG. 48C. Client 4910 may repeat any of steps 4912 and 4914, and server 4920 may repeat step 4922, as many times as necessary to transmit data between client 4910 and server 4920. Secure proxy service 4900 then ends.

FIG. 50 is a simplified block diagram of a secure proxy service 5000 between client 5010 and server 5020 that distributes trust among a set of certificate authorities in the structure of communication channels, that may be used in any suitable combination, with any suitable additions, deletions, or modifications in accordance with one embodiment of the present invention. Secure proxy service 5000 may be, for example, the secure proxy services discussed with respect to process flow 4700, 4800, or 4850, as well as secure proxy service 4870 or 4900 as described with respect to FIGS. 47-49. Secure proxy service 5000 may be implemented such that each one of the secure communication tunnels 5030, 5040, and 5050 is associated with an outer level certificate authority (not shown in FIG. 50) as described with respect to step 4710 of process flow 4700 of FIG. 47 and step 4810 of process flow 4800 of FIG. 48A. In addition, secure proxy service 5000 may be implemented such that each one of secure communication tunnels 5030, 5040, and 5050 is associated with a different one of an inner level certificate authority such as first certificate authority CA1, second certificate authority CA2, and third certificate authority CA3, and each one of the secure communication tunnels 5030, 5040, and 5050 are established over different physical transport mediums. These different physical media may be any suitable physical transport medium as described with respect to step 4710 of process flow 4700 of FIG. 47 or with respect to step 4810 of process flow 4800 of FIG. 48A. For example, as illustrated in FIG. 50, secure communication tunnels 5030,

5040, and 5050 may be established over WiFi, Ethernet, and cellular communication channels, respectively.

During normal operation of secure proxy service 5000, client 5010 may send first data share 5012 over first secure communication tunnel 5030, second data share 5014 over second secure communication tunnel 5040, and third data share 5016 over third secure communication tunnel 5050. Data shares 5012, 5014, and 5016 may be data shares computed by client 5010 using any suitable data dispersing and encryption techniques discussed with respect to step 4730 of process flow 4700 of FIG. 47, or discussed with respect to step 4820 and 4825 of process flow 4800 of FIG. 48A. Although secure proxy service 5000 is illustrated as splitting data into 3 data shares, it may be understood that secure proxy service 5000 may disperse data into any suitable number of shares and transmit them each over any suitable number of secure communication tunnels.

In some embodiments, one of the physical media may experience a network failure. This network failure may be due to a structural malfunction of the physical media. For example, as illustrated in secure proxy service 5000, the cellular communication channel used to establish third secure communication tunnel 5050 may experience a network failure due to damage to a cellular tower.

In some embodiments, secure proxy service 5000 may not change the transmission of its data packets over secure communication tunnels 5030, 5040, and 5050 in response to the network failure (not illustrated in FIG. 50). In other words, client 5010 may continue to send first data share 5012 over first secure communication tunnel 5030, and second data share 5014 over second secure communication tunnel 5040. In some embodiments of secure proxy service 5000, the data packets that include first data share 5012 and second data share 5014 may continue to be exchanged without a loss of data integrity. For example, if the data packets processed by secure proxy service 5000 are split using a secret sharing algorithm (e.g., the Shamir secret sharing algorithm) such that only a quorum of data shares 5012, 5014, and 5016 are needed to recover each transmitted data packet, then there may be no loss of data integrity between the packets exchanged between client 5010 and server 5020.

In some embodiments, secure proxy service 5000 may change the transmission of its data packets over secure communication tunnels. For example, first data share 5012 may continue to be transmitted over the wifi-based secure communication tunnel 5030, while second data share 5014 and third data share 5014 may be transmitted over the Ethernet-based secure communication tunnel 5040. In such embodiments, one or more of the secure communication tunnels may need to be split into one or more secure communication tunnels using additional key establishment processes. This split may be accomplished according to, for example, any of the key establishment phases described with respect to step 4720 of process flow 4700 of FIG. 47 or step 4820 of process flow 4800 of FIG. 48. After these additional key establishment phases, the data shares of the data packets may resume according to the new configuration of the secure proxy service 5000.

In some embodiments, this new configuration of secure proxy service 5000 may change how the split shares of data are encrypted based on keys associated with the establishment of the secure communication tunnels or split portions of the secure communication tunnels. For example first data share 5012 may be encrypted based on a key associated with the establishment of wifi-based first secure communication tunnel 5030, and then transmitted over that tunnel. Second data share 5014 may be encrypted based on a key associated

with the establishment of a first split portion 5042 of Ethernet-based second secure communication tunnel 5040 and transmitted over first split portion 5042, and third data share 5016 may be encrypted based on a key associated with the establishment of a second split portion 5044 of Ethernet-based secure communication tunnel 5040 and transmitted over second split portion 5044. In some embodiments, first secure communication tunnel 5030 may be established using the certificate obtained from certificate authority CA1, first split portion 5042 of Ethernet-based second secure communication tunnel 5040 may be established using the certificate obtained from certificate authority CA2, and second split portion 5044 of Ethernet-based secure communication tunnel 5040 may be established using the certificate obtained from certificate authority CA3. In some embodiments, secure proxy service 5000 may execute these additional key establishment processes adaptively as the communication channels that the secure communication tunnels are established on fail or are restored. In embodiments of secure proxy service 5000 as illustrated in FIG. 50, secure proxy service 5000 may be referred to as “communication-medium redundant”.

In embodiments of secure proxy service 5000 as illustrated in FIG. 50, data may be exchanged between client 5010 and server 5020 without a loss of data confidentiality, integrity, and authenticity. That is, even if compromised certificate authorities have access to the information exchanged based on that certificate authority, the attacker associated with that certificate authority may not have enough information to disrupt the confidentiality or integrity of the communication. For example, if the outer level certificate authority of secure proxy service 5000 was compromised but the integrity of the inner-level certificate authorities were preserved, the attacker may be able to observe the streams of data over the WiFi-based communication channel and both portions of the Ethernet-based communication channel, but may have no knowledge of the encryption used to secure data over each secure communication tunnel.

In another example, if one or more of the inner level certificate authorities was compromised but the outer level certificate authority remained intact, the attacker may be able to recover some of the cryptographically split portions of the data packets, but may not be able to decrypt the data packets themselves because it may have no knowledge of the encryption used by the outer level certificate authority. Further, if the data packets sent through the secure communication tunnels are cryptographically split such that they are restorable from at least a subset of the shares by recombining at least a quorum of the shares, the user of the secure data parser may have the additional protection that if some, but less than a quorum, of the certificates associated with the secure communication tunnels have been compromised, the attacker may not be able to restore the cryptosplit data packets.

Additionally, other combinations, additions, substitutions and modifications will be apparent to the skilled artisan in view of the disclosure herein.

What is claimed is:

1. A method for securing data in motion comprising original data packets, the method comprising:
 - establishing a secure communication channel;
 - establishing a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities;
 - dispersing each one of the original data packets into a plurality of shares based on multi-factored secret sharing;

107

encrypting each of the plurality of shares using a key associated with the establishment of a different one of the secure communication tunnels; and
 transmitting the plurality of encrypted shares over one or more of the plurality of secure communication tunnels.

2. The method of claim 1, wherein transmitting the plurality of encrypted shares comprises transmitting each encrypted share on a secure communication tunnel established using a key used to encrypt the respective encrypted share.

3. The method of claim 1, further comprising:
 receiving the plurality of encrypted shares, each on a respective one of the plurality of secure communication tunnels;
 decrypting each of the plurality of encrypted shares based on the key associated with the establishment of the respective one of the plurality of secure communication tunnels; and
 restoring the set of original data packets based on the multi-factored secret sharing.

4. The method of claim 1, further comprising generating a certificate authority hierarchy, wherein the certificate authority hierarchy comprises a set of root certificate authorities, and wherein the plurality of unique certificate authorities comprises the set of root certificate authorities.

5. The method of claim 1, further comprising generating a certificate authority hierarchy, wherein the certificate authority hierarchy comprises a set of minor certificate authorities, and wherein the plurality of unique certificate authorities comprises the set of minor certificate authorities.

6. The method of claim 1, wherein each of the plurality of secure communication tunnels is established over a different physical transport medium.

7. The method of claim 6, wherein at least one of the physical transport mediums experiences a network failure, the method further comprising restoring the original data packets without a loss of data integrity.

8. The method of claim 7, wherein at least one of physical transport mediums experiences a network failure, wherein a portion of the plurality of shares is designated for transmission over the failed physical transport medium, and wherein at least one of the physical transport mediums is operational, the method further comprising:
 establishing additional secure communication tunnels within the at least one operational physical transport medium; and
 transmitting the portion of the plurality of shares designated for transmission over the at least one of the set of failed physical transport mediums over the additional secure communication tunnels.

9. The method of claim 1, wherein the associations between the tunnels and the unique certificate authorities are dynamic.

10. The method of claim 1, wherein dispersing each one of the original data packets into a plurality of shares further comprises splitting each one of the original data packets into a plurality of shares based on an M of N cryptosplit.

11. The method of claim 10, wherein the plurality of shares are restorable from at least a subset of the shares by recombining at least a quorum of the shares.

12. The method of claim 1, wherein the plurality of secure communication tunnels are established based on the Transport Layer Security protocol.

13. The method of claim 1, further comprising:
 generating a certificate authority hierarchy, wherein the certificate authority hierarchy comprises a set of minor certificate authorities; and

108

encrypting each one of the set of shares based on a certificate issued by a unique minor certificate authority of the certificate authority hierarchy.

14. The method of claim 1, wherein transmitting the plurality of encrypted shares comprises transmitting two or more of the plurality of shares over a same one of the plurality of secure communication tunnels.

15. The method of claim 1, wherein transmitting the plurality of encrypted shares comprises transmitting each of the plurality of shares over one of the plurality of secure communication tunnels, wherein each share is transmitted over a secure communication tunnel that is different than the tunnel corresponding to the key used to encrypt the respective share.

16. The method of claim 1, wherein transmitting the plurality of encrypted shares comprises transmitting each of the plurality of encrypted shares over a different one of the plurality of secure communication tunnels.

17. The method of claim 1, wherein transmitting the plurality of encrypted shares comprises transmitting the plurality of shares over the plurality of secure communication tunnels, wherein at least two of the plurality of encrypted shares are transmitted over the same tunnel.

18. The method of claim 1, wherein establishing a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities comprises establishing the plurality of secure communication tunnels using a different one of the certificate authorities to establish each of the secure communication tunnels.

19. The method of claim 1, wherein establishing a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities comprises establishing the plurality of secure communication tunnels using a same one of the certificate to establish at least two, but less than all, of the plurality of secure communication tunnels.

20. A system for securing data in motion comprising original data packets, the system comprising a first device comprising processing circuitry configured to:
 establish a secure communication channel;
 establish a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities;
 disperse each one of the original data packets into a plurality of shares based on multi-factored secret sharing;
 encrypt each of the plurality of shares using a key associated with the establishment of a different one of the secure communication tunnels; and
 transmit the plurality of encrypted shares over one or more of the plurality of secure communication tunnels.

21. The system of claim 20, wherein the processing circuitry is configured to transmit the plurality of encrypted shares by transmitting each encrypted share on a secure communication tunnel established using a key used to encrypt the respective share.

22. The system of claim 20, further comprising a second device comprising processing circuitry further configured to:
 receive the plurality of encrypted shares, each on a respective one of the plurality of secure communication tunnels;

109

decrypt each of the plurality of encrypted shares based on the key associated with the establishment of the respective one of the plurality of secure communication tunnels; and

restore the set of original data packets based on the multi-
factored secret sharing.

23. The system of claim 20, wherein the processing circuitry is further configured to generate a certificate authority hierarchy, wherein the certificate authority hierarchy comprises a set of root certificate authorities, and wherein the plurality of unique certificate authorities comprises the set of root certificate authorities.

24. The system of claim 20, wherein the processing circuitry is further configured to generate a certificate authority hierarchy, wherein the certificate authority hierarchy comprises a set of minor certificate authorities, and wherein the plurality of unique certificate authorities comprises the set of minor certificate authorities.

25. The system of claim 20, wherein each of the plurality of secure communication tunnels is established over a different physical transport medium.

26. The system of claim 25, wherein at least one of the physical transport mediums experiences a network failure, the method further comprising restoring the original data packets without a loss of data integrity.

27. The system of claim 26, wherein at least one of physical transport mediums experiences a network failure, wherein a portion of the plurality of shares is designated for transmission over the failed physical transport medium, and wherein at least one of the physical transport mediums is operational, the processing circuitry further configured to:

establish additional secure communication tunnels within the at least one operational physical transport medium; and

transmit the portion of the plurality of shares designated for transmission over the at least one of the set of failed physical transport mediums over the additional secure communication tunnels.

28. The system of claim 20, wherein the associations between the tunnels and the unique certificate authorities are dynamic.

29. The system of claim 20, wherein the processing circuitry is further configured to disperse each one of the original data packets into a plurality of shares by splitting each one of the original data packets into a plurality of shares based on an M of N cryptosplit.

110

30. The system of claim 29, wherein the plurality of shares are restorable from at least a subset of the shares by recombining at least a quorum of the shares.

31. The system of claim 20, wherein the plurality of secure communication tunnels are established based on the Transport Layer Security protocol.

32. The system of claim 20, wherein the processing circuitry is further configured to transmit the plurality of encrypted shares by transmitting the plurality of shares over one of the plurality of secure communication tunnels.

33. The system of claim 20, wherein the processing circuitry is further configured to transmit the plurality of encrypted shares by transmitting each of the plurality of shares over one of the plurality of secure communication tunnels, wherein each share is transmitted over a secure communication tunnel that is different than the tunnel corresponding to the key used to encrypt the respective share.

34. The system of claim 20, wherein the processing circuitry is further configured to transmit the plurality of encrypted shares by transmitting each of the plurality of shares over a different one of the plurality of secure communication tunnels.

35. The system of claim 20, wherein the processing circuitry is further configured to transmit the plurality of encrypted shares by transmitting the plurality of shares over the plurality of secure communication tunnels, wherein at least two of the plurality of encrypted shares are transmitted over the same tunnel.

36. The system of claim 20, wherein the processing circuitry is further configured to establish a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities by establishing the plurality of secure communication tunnels using a different one of the certificate authorities to establish each of the secure communication tunnels.

37. The system of claim 20, wherein the processing circuitry is further configured to establish a plurality of secure communication tunnels within the secure communication channel, wherein the plurality of secure communication tunnels is established using certificates issued by a plurality of unique certificate authorities by establishing the plurality of secure communication tunnels using a same one of the certificate to establish at least two, but less than all, of the plurality of secure communication tunnels.

* * * * *